Fictitious Play with Inertia Learns Pure Equilibria in Distributed Games with Incomplete Information

Brian Swenson[1], Ceyhun Eksin[2], Soummya Kar[1], Alejandro Ribeiro[3]

arXiv:1605.00601v2 [cs.SY] 6 Jan 2017

***Abstract*—The paper studies a class of learning dynamics, termed *inertial best response dynamics*, based on the Fictitious Play (FP) algorithm. A general convergence result is established showing that inertial best response dynamics converge, almost surely, to pure-strategy Nash equilibrium (NE). As an application of the general result, the paper considers learning in a setting where (i) players have some uncertainty about the underlying state of the world, and (ii) all inter-agent communication is restricted to a preassigned (possibly sparse) communication graph. The paper studies two particular instances of inertial best response dynamics in this setting: FP with inertia and Joint Strategy FP with inertia. General conditions are established under which the algorithms converge to pure-strategy NE, and a fully distributed variant of each algorithm is presented. Finally, numerical simulations are provided which verify the findings.**

I. Introduction

Game theoretic learning processes have gained increasing prominence in the control, signal processing, and communication communities due to their applicability as decentralized algorithms [1]–[6]. In these contexts, a multi-agent system is often modeled using the framework of a normal-form game [1], [7]. Pure-strategy Nash equilibria of the game can often be shown to represent desirable operating conditions in the associated multi-agent system (e.g. [8]).

There exist many game-theoretic learning algorithms that ensure agents learn pure-strategy NE strategies (e.g., [9], [10]). However, such algorithms generally assume that agents have instantaneous access to information necessary to compute their next-stage action. In many real-world scenarios, global information is not instantly available to all agents—rather, information tends to be distributed throughout the system and must be disseminated to all agents using some overlaid communication infrastructure, e.g., wireless networks [2], [4], multi-robot systems [8], smart grid infrastructures [5], [11].

In this paper, we consider a setting in which each agent has access to private information about their own action history, but does not have (direct) access to information about the actions of others, nor access to instantaneous payoff information. We assume that agents are endowed with an overlaid communication graph where a vertex represents a player and an edge denotes the ability for two players to exchange information (e.g., see [12]–[16]). An agent may only obtain information related to another agent's action history if that information is disseminated through the communication graph.

The assumption that players do not have access to payoff information is relevant in settings where feedback from physical game play is not available, or is delayed. This may occur, for example, in scenarios where agents wish to precompute an equilibrium strategy prior to physical game play, or where the designed utility functions do not readily admit physical measurement. In such a framework, information can be subject to corruption and delay as it is disseminated through the communication graph, and the traditional convergence results of pure-strategy-NE seeking algorithms do not apply.

In addition to assuming that information flow is restricted to an inter-agent communication graph, we suppose that agents may have some uncertainty about the state of the world. We model this uncertainty using a game of incomplete information in which players have some uncertainty about the game utility functions. We assume that each player receives a private signal that is correlated to the true (time-invariant) state of the world. Using this information, each agent forms a personal belief about the state of the world, which they may update as information is shared over the communication graph through the course of the learning process. A learning process in this setting must take into account the joint objectives of uncertainty mitigation and pure-strategy NE learning.

We consider a broad class of learning dynamics based on the well-known Fictitious Play (FP) algorithm [17], [18] and demonstrate convergence results within the class of weakly acyclic games [19].

The problem of learning Nash equilibria (possibly mixed) through distributed implementations of FP-type algorithms in multi-agent networks has been addressed in prior work [12], [15], [16]. In this paper we focus on distributed equilibria learning in which the agents specifically seek to learn *pure-strategy* Nash equilibria which is not guaranteed by the distributed learning mechanisms presented in [12], [15] and also, in settings with incomplete information [16].

In order to learn pure-strategy equilibria in this setting we consider a class of algorithms we denote as *inertial best response dynamics*— such algorithms consist of a best-response component (similar to classical FP) and an additional "inertial term" which ensures that players occasionally repeat actions in consecutive stages. The use of an inertial term is a common technique used to ensure convergence to pure-strategy NE [9], [10].

We first prove a general convergence result that applies to

*The work of B. Swenson and S. Kar was supported in part by NSF grant CCF-1513936.

[1]Dept. of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA. brianswe@ece.cmu.edu, soummyak@andrew.cmu.edu

[2]School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332. ceyhuneksin@gatech.edu

[4]Dept. of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA. aribeiro@seas.upenn.edu

any inertial best response algorithm. The general convergence result applies to a broad class of learning dynamics and applies to a wide variety of information dissemination schemes (not just the synchronous graph-based information dissemination schemes that we focus on in the later part of the paper).

Subsequently, as applications of the general result, we study the distributed implementation (under uncertainty) of two important inertial best response algorithms.

As a first application, we study the distributed implementation of classical FP with inertia (referred to as D-FP) (see Section IV).

As a second application, we study the distributed implementation of a variant of the Joint Strategy FP algorithm [9] that is applicable within the class of congestion games. (We refer to this distributed algorithm as D-JSFP.) Classical FP can be difficult to implement in large-scale games due to high demands in terms of information overhead and complexity as the number of players grows large. JSFP is a variant of FP that mitigates these issues. We note that while the D-JSFP algorithm of this paper is able to operate with lower computational and communication burdens, it is applicable within a narrower class of games than D-FP. (This is a consequence of the structure of the information that is passed through the network—see Section V for more details.)

In summary, the paper investigates the problem of learning pure-strategy NE in games where there is environmental uncertainty and players must communicate information using an overlaid communication graph. As our first main contribution, we study a general class of learning algorithms that we call inertial best response dynamics, and we prove such algorithms converge to pure-strategy NE under appropriate assumptions. Our second main contribution is an application of this general result to develop variants of the FP and JSFP algorithms that achieve pure-strategy NE learning in distributed settings with uncertainty.

The remainder of the paper is organized as follows. In Section II we set up the problem formulation. In Section III we present the general convergence result for inertial best response dynamics. In Section IV we study distributed implementations of classical FP with inertia. In Section V we study distributed implementations of the Joint Strategy FP algorithm with inertia. In Section VI we present simulation examples.

II. Networked multi-agent systems with incomplete information

In this paper we consider scenarios where players in a game have incomplete information about the state of the world. Formally, we begin by defining Θ to be the set of possible states of the world and we assume that Θ is a metrizable space. We let $\mathcal{B}(\Theta)$ denote the Borel σ-algebra on Θ and let $\Delta(\Theta)$ denote the set of probability measures on the measurable space $(\Theta, \mathcal{B}(\Theta))$. Further use $\mathcal{N} = \{1, \dots, n\}$ to denote the set of players and $\mathcal{A}_i$ to denote the action space of player i, which we assume to be finite. The joint action space of all the players is then given by the Cartesian product $\mathcal{A} = \prod_{i=1}^n \mathcal{A}_i$.

For each state of the world $\theta \in \Theta$ we consider the associated normal-form game defined by the tuple

$$\Gamma(\theta) = (\mathcal{N}, \{\mathcal{A}_i, u_i(\cdot, \theta)\}_{i \in \mathcal{N}}), \tag{1}$$

where $u_i(\cdot, \theta) : \mathcal{A} \to \mathbb{R}$ denotes the utility function of player i given the state θ. Observe that the set of players and the set of possible actions are assumed to be the same irrespective of the state of the world θ. When player i chooses strategy a_i, the payoff it receives is $u_i(\{a_i\}_{i \in \mathcal{N}}, \theta) = u_i(a, \theta)$. The payoff depends on the joint action profile $a = \{a_i\}_{i \in \mathcal{N}}$ and the state of the world θ. We suppose that players are permitted to use probabilistic strategies. Formally, let the *mixed strategy space* of player i be given by the set $\Delta(\mathcal{A}_i)$ of probability distributions over $\mathcal{A}_i$ and let $\Delta^n(A) := \prod_{i=1}^n \Delta(\mathcal{A}_i)$ denote the set of joint mixed strategies where it is assumed that players use independent strategies. We represent a joint mixed strategy $\sigma \in \Delta^n(\mathcal{A})$ as the n-tuple $\sigma = \{\sigma_i\}_{i \in \mathcal{N}}$, where $\sigma_i \in \Delta(\mathcal{A}_i)$ denotes the marginal mixed strategy of player i. When a mixed strategy $\sigma \in \Delta^n(\mathcal{A})$ is played, we are interested in the *expected* payoffs which we write as

$$u_i(\sigma, \theta) \;=\; u_i(\sigma_i, \sigma_{-i}, \theta) \;=\; \sum_{a \in \mathcal{A}} u_i(a, \theta)\sigma(a). \tag{2}$$

The notation $u_i(\sigma_i, \sigma_{-i}, \theta)$ in (2) is meant to emphasize that the payoff depends on the strategy σ_i chosen by player i and the strategies $\sigma_{-i} := \{\sigma_j\}_{j \in \mathcal{N}, j \neq i}$ that are chosen by other players.

Since we want to study games in which θ is uncertain, we further define expected payoffs with respect to a probability distribution $\mu \in \Delta(\Theta)$—we refer to μ as a *belief* about the state of the world. The expectation of the payoff in (2) with respect to this belief is then given by

$$u_i(\sigma, \mu) \;=\; u_i(\sigma_i, \sigma_{-i}, \mu) \;=\; \int u_i(\sigma, \theta)\mu(d\theta). \tag{3}$$

The notation $u_i(\sigma, \mu)$ in (3) is a slight abuse of the notation $u_i(\sigma, \theta)$ in (2) which is itself an slight abuse of $u_i(a, \theta)$. These inconsistencies are resolved if we use θ to denote a state realization but also as a shorthand for the belief μ with all mass on θ, and use a_i and a to denote actions but also as a shorthand for strategies σ_i and σ with all mass in a_i and a. We adopt this convention henceforth.

The payoffs in (3) define a modification of the game $\Gamma(\theta)$ in (1) in which the payoffs are given by the expected utilities in (3). The equilibria of this modified game are formally defined next.

Definition 1 *Given the game $\Gamma(\theta)$ in* (1) *and a belief μ define the game $\Gamma(\mu) = (\mathcal{N}, \{\mathcal{A}_i, u_i(\cdot, \mu)\})$ associated with the utilities in* (3)*. The joint strategy $\sigma = \{\sigma_i\}_{i \in \mathcal{N}} \in \Delta^n(\mathcal{A})$ is a Nash equilibrium of $\Gamma(\mu)$ if the utilities in* (3) *satisfy*

$$u_i(\sigma_i, \sigma_{-i}, \mu) \geq u_i(\sigma_i', \sigma_{-i}, \mu), \quad \text{for all } \sigma_i' \in \Delta(\mathcal{A}_i),\ i \in \mathcal{N}. \tag{4}$$

The equilibrium is said to be a pure-strategy equilibrium if there exists an action tuple $a \in \mathcal{A}$ such that $\sigma = a$.

The equilibria in (4) exist because, with μ fixed, the payoffs in (3) define a normal form game. In general, pure equilibria

may or may not exist, but we will assume here that they do (see Assumption 2). In such case, identifying pure equilibria is important in applications because they result in behaviors that are often more reasonable than the behaviors that result from mixed equilibria—see Section I. Our goal in this paper is to design an algorithm allowing agents to learn the pure equilibria in Definition 1 when players play repeatedly over time as they acquire a common belief on the state of the world. We formally state this problem in the following section.

A. Learning via repeated play

We suppose that players are permitted to repeatedly face off in the game $\Gamma(\theta)$ in discrete stages $t \in \mathbb{N}_+$. At each stage t, agent i plays an action $a_{i,t} \in \mathcal{A}_i$ that is chosen from the possibly mixed strategy $\sigma_{i,t} \in \Delta(\mathcal{A}_i)$. To select the strategy $\sigma_{i,t}$, agent i has access to a belief $\mu_{i,t}$ on the state of the world, and has some knowledge of the past history of game play $\{a_{j,s}\}_{s=1}^{t-1}$, $j \in \mathcal{N}$. The beliefs $\mu_{i,t} \in \Delta(\Theta)$ are time varying and possibly different for different agents, although they are required to converge to a common distribution—see Assumption 1. Besides this belief on the state of the world, agent i also has estimates $\hat{\sigma}^i_{j,t} \in \Delta(\mathcal{A}_j)$ of the strategies $\sigma_{j,t} \in \Delta(\mathcal{A}_j)$ of other agents. The belief $\mu_{i,t}$ and the estimated strategies $\hat{\sigma}^i_{-i,t} := \{\hat{\sigma}_{j,t}\}_{j\in\mathcal{N}, j\neq i}$ allow agent i to estimate the payoff that it would receive from playing an arbitrary action a_i. These estimated payoffs can be written as [cf. (3)]

$$\hat{u}_{i,t}(a_i) := u(\sigma_i, \hat{\sigma}^i_{-i,t}, \mu_{i,t}), \quad a_i \in \mathcal{A}_i. \tag{5}$$

The utility $\hat{u}_{i,t}(a_i)$ in (5) is what agent i predicts he would receive as payoff given the (partial) information he has available on the state of the world (the belief $\mu_{i,t}$) and the actions of other agents (the strategy estimates $\hat{\sigma}^i_{-i,t}$).

The action that maximizes the utility estimates in (5) is of importance to agent i and defined here as the best response

$$\hat{a}_{i,t} \;:=\; \operatorname*{argmax}_{a_i\in\mathcal{A}_i} \hat{u}_{i,t}(a_i) \;=\; \operatorname*{argmax}_{a_i\in\mathcal{A}_i} u(a_i, \hat{\sigma}^i_{-i,t}, \mu_{i,t}). \tag{6}$$

Observe that there may be multiple arguments that maximize $\hat{u}_{i,t}(a_i)$. Thus, the best response $\hat{a}_{i,t}$ in (6) is in general a set, not an individual action.

A learning algorithm in game theory is a collection of behavior rules that dictate how each player should choose the "next-stage action" $a_{i,t+1}$ based on the information they have available at time t. Myopic best response rules in which agents play an action $a_{i,t+1} \in \hat{a}_{i,t}$ are common; see e.g., [7], [10], [17]. Here, we are interested in a slight modification in which agents are sometimes "reluctant" to modify their action choices from round to round. We refer to this general algorithm (formally stated next) as *inertial best response dynamics*.[1]

Algorithm 1 Let $\rho \in (0,1)$ be an inertia constant and let $a_{i,1}$ be an arbitrary initial action for each i. At time $t > 1$, agent i has access to strategy estimates $\hat{\sigma}^i_{j,t}$ and a belief $\mu_{i,t}$ that it uses to compute the best response set $\hat{a}_{i,t}$ in (6). Players are said to follow inertial best response dynamics if they play actions according to

$$\mathbb{P}\left(a_{i,t+1} = a_{i,t} | \mathcal{F}_{t-1}\right) = \rho, \quad \mathbb{P}\left(a_{i,t+1} \in \hat{a}_{i,t} | \mathcal{F}_{t-1}\right) = 1-\rho, \tag{7}$$

where $(\mathcal{F}_t)_{t\geq 1}$ is a filtration (sequence of increasing σ-algebras) that contains the information available to players in round t. As per (7), an inertial best response algorithm entails player i sticking to its previous play $a_{i,t}$ with some (fixed) probability ρ (this is the inertia component of the algorithm), and playing a best response otherwise. The selection of myopic best responses promotes a general trend of improving utility, while the inertia component of the algorithm—with its occasional repetition of actions—ensures that players eventually lock into a pure strategy equilibrium. Examples of algorithms that use best response with inertia can be found in [9], [10].

Throughout this paper, we want to study mechanisms to update the beliefs $\hat{\sigma}^i_{j,t}$ so that players learn pure equilibria (cf. Definition 1). We do so using ideas from fictitious play, so that (roughly speaking) $\hat{\sigma}^i_{j,t}$ is built from empirical histograms of the actions that have been observed or action information that has been disseminated through an overlaid communication graph. These algorithms are defined in Sections IV and V after we introduce some preliminary convergence results for the generic inertial best response methods given in Algorithm 1.

III. Inertial Best Response Algorithm: General Convergence Result

In this section we study convergence results for the inertial best response learning dynamics presented in the previous section (see Algorithm 1). These dynamics can be shown to achieve pure-strategy NE learning so long as the estimated payoffs $\hat{u}_{i,t}$ satisfy an appropriate asymptotic accuracy condition (see Condition 1). We first present some general assumptions regarding game play, and then present the convergence results.

In later sections, these results will be used to prove convergence of two particular instances of inertial best response dynamics.

A. General Assumptions

We first introduce some notation necessary for the subsequent development. Suppose S is a metric space with Borel σ-algebra Σ, and let $\{\mu_t\}_{t\geq 1}$ and μ be probability measures on (S, Σ). The sequence $\{\mu_t\}_{t\geq 1}$ converges weakly to μ if $\int g(s)\mu_t(ds) \to \int g(s)\mu(ds)$ as $t \to \infty$ for any bounded continuous function $g : S \to \mathbb{R}$. In this case we write $\mu_t \xrightarrow{w} \mu$.

We assume that players' beliefs on the state of the world converge weakly to a fixed common distribution. Formally,

Assumption 1 *There exists a probability measure μ such that $\mu^i_t \xrightarrow{w} \mu$ for all $i \in \mathcal{N}$.*

This assumption is easily attainable by a variety of belief formation mechanisms. For example, suppose that there exist sequences $(\theta^i_t)_{t\geq 1} \subset \Theta$ satisfying $\lim_{t\to\infty} \theta^i_t = \theta$, for all i.

[1] We also refer to an algorithm of this form as an *inertial best response algorithm.*

Then the belief sequence where μ_t^i places mass 1 on θ_t^i satisfies Assumption 1.

Throughout the paper, our focus is aimed at studying learning dynamics and not the particular process used to form the beliefs μ_t^i. In an effort to maintain this focus, we will not explicitly specify any particular dynamics used to form the belief μ_t^i—rather, we will study learning results that can be achieved under the broad assumption of weak convergence of beliefs. In Section VI we discuss a simulation example in which the dynamics used to form the belief μ_t^i are explicitly specified and satisfy Assumption 1.

Given that there exists a μ such that $\mu_t^i \xrightarrow{w} \mu$, for all i, we assume further that under the belief μ, the game is weakly acyclic. Formally,

Assumption 2 *Let μ be as in Assumption 1 and define the normal form game Γ associated with the belief μ as $\Gamma(\mu) = (\mathcal{N}, \{\mathcal{A}_i, u_i(\cdot, \mu)\})$. We assume the game $\Gamma(\mu)$ is weakly acyclic; that is, for any $a \in \mathcal{A}$, there exists a best-response path that converges to a pure-strategy Nash equilibrium.*

A review of the topic of weakly acyclic games can be found in [19]. The above assumption ensures that the set of pure-strategy NE (see Definition 1) is non-empty, and will ensure that the learning algorithms to be studied can always reach a pure-strategy NE.

Finally, we assume that under the limiting belief μ, players are never indifferent between action choices. Formally,

Assumption 3 *For any pair of actions $a_i' \in \mathcal{A}_i$, $a_i'' \in \mathcal{A}_i$, $a_i' \neq a_i''$ and for any $a_{-i} \in \mathcal{A}_{-i}$,*

$$u_i(a_i', a_{-i}, \mu) \neq u_i(a_i'', a_{-i}, \mu)$$

for all $i \in \mathcal{N}$.

Note that Assumption 3 (cf. [9] Assumption 2.2) implies that under the belief μ all pure NE are strict. Moreover, the assumption implies that for any $i \in \mathcal{N}$ and $a_{-i} \in \mathcal{A}_{-i}$, the set $\operatorname{argmax}_{\alpha_i \in \mathcal{A}_i} u_i(\alpha_i, a_{-i}, \mu)$ is a singleton.

We remark that Assumption 3 is generic in the sense that if the number of players and actions are fixed, then the set of utility functions for which Assumption 3 fails is a closed set of Lebesgue measure zero within the space of all possible utility functions.

Assumption 4 *Let $\{\mathcal{F}_t\}_{t\geq 1}$ be a filtration (sequence of increasing σ-algebras) with $\mathcal{F}_t := \sigma(\{a_s\}_{s=1}^t, \{\mu_s^i\}_{i\in\mathcal{N}, s=1,\ldots,t})$. The strategy estimate $\hat{\sigma}_{j,t}^i \in \Delta(\mathcal{A}_j)$ that agent i has of the strategy $\sigma_{j,t}$ of agent j is measurable with respect to $\mathcal{F}_t$.*

Assumption 4 means that the strategy estimates of agent i are restricted to be a function of the history of play.

The following condition is sufficient (to be shown) to ensure that a FP-type algorithm with inertia converges to NE.

Condition 1 *There exist $\bar{t}$, $T \in \mathbb{N}_+$ such that, if $t \geq \bar{t}$, and if any action $a \in \mathcal{A}$ is repeated consecutively for $\tilde{T} \geq T$ stages (i.e., $a_s = a$ for $s = t, \ldots, t + \tilde{T} - 1$), then $\operatorname{argmax}_{\alpha_i \in \mathcal{A}_i} \hat{u}_{i,t+\tilde{T}-1}(\alpha_i) = \operatorname{argmax}_{\alpha_i \in \mathcal{A}_i} u_i(\alpha_i, a_{-i}, \mu)$ for all i.*

Throughout this section we assume that this condition holds. (In later sections, when studying particular inertial best response algorithms, our main task will be to show that Condition 1 holds.) The condition states that there exists a finite time $\bar{t}$ after which if any action profile a is repeated a sufficiently large number of times (specifically, more than T times) then player i's optimal response to its predictions of the utility coincides with the optimal response to the actual action profile a_{-i}. This condition is natural within the class of learning algorithms with fading or finite memory [10] where, if players repeat an action a sufficient number of times their prediction of the utility can be brought arbitrarily close to the utility of the repeatedly-played action.

When studying particular implementations of inertial best response dynamics in the following sections, the condition reduces to tangible assumptions on learning past action profiles of other players—see Conditions 2 and 3.

B. General Convergence Analysis

The following two lemmas allow us to prove the main theoretical result of this section. Lemma 1 shows that pure-strategy Nash equilibria are absorbing, and Lemma 2 shows that the probability of reaching such an absorbing state is uniformly bounded from below. Together they prove Theorem 1.

Lemma 1 (absorption property) *Let $\bar{t}$ be as in Condition 1. Let $\{a_t\}_{t\geq 1}$ be a sequence of actions generated by an inertial best response algorithm. Suppose Assumptions 1–4 and Condition 1 hold. There exists a $T_1 \in \mathbb{N}_+$ such that if $t \geq \bar{t}$, and $a^* \in \mathcal{A}$ is any pure-strategy Nash equilibrium, and if a^* is played in T_1 consecutive stages, i.e., $a_s = a^*$, for all $s = t, \ldots, t + T_1 - 1$, then $a_{t+\tau} = a^*$ for all $\tau \geq 0$.*

Proof : Let T be as in Condition 1, and let $T_1 \geq T$. Suppose a^* is a pure Nash equilibrium and $a_s = a^*$ for $s = t, \ldots, t + T_1 - 1$. Then by Condition 1, $\operatorname{argmax}_{a_i' \in \mathcal{A}_i} \hat{u}_{i,t+\tilde{T}}(a_i') = \operatorname{argmax}_{a_i' \in \mathcal{A}_i} u_i(a_i', a_{-i}^*, \mu)$, for all i (recall that $\hat{u}_{i,t+\tilde{T}}(a_i')$ as defined in Section II-A is implicitly dependent on $\mu_{t+\tilde{T}}^i$). Moreover, by Assumption 3, the set $\operatorname{argmax}_{a_i' \in \mathcal{A}_i} u_i(a_i', a_{-i}^*, \mu) = \{a_i^*\}$ is a singleton for each i. Thus, the action a^* is repeated in stage $t + T_1$. Inductively, we see that $a_{t+\tau} = a^*$ for all $\tau \geq 0$. □

Lemma 2 (positive probability of absorption) *Let $\bar{t}$ be as in Condition 1. Let $\{a_t\}_{t\geq 1}$ be a sequence of actions generated by an inertial best response algorithm. Suppose Assumptions 1–4 and Condition 1 hold. Let T_1 be as in Lemma 1, and let $T_2 \geq T_1$ be given. Let t be the current stage of the repeated play. Define the event*

$$E_t := \{a_\tau = a^* \text{ for some pure strategy NE } a^* \text{ for all } \tau \in \{t', t'+1, \ldots, t'+T_2-1\}, \text{ for some } t' \in \{t, \ldots, t + T_2|\mathcal{A}|\}\}.$$

There exists an $\epsilon = \epsilon(T_2) > 0$ such that $\mathbb{P}(E_t|\mathcal{F}_t) > \epsilon$ for all $t \geq \bar{t}$.

Proof : The proof follows along the lines of the proof of Theorem 3.1 in [9]. By Condition 1, for any $t' \geq \bar{t}$, if any action $a \in \mathcal{A}$ is repeated consecutively from stage t' to stage $t' + T_2 - 1$, then $\text{argmax}_{\alpha_i \in \mathcal{A}_i} \hat{u}_{i,t'+T_2-1}(\alpha_i) = \text{argmax}_{\alpha_i \in \mathcal{A}_i} u_i(\alpha_i, a_{-i}, \mu)$. Let $a^0 = a_t$. Conditioned on $\mathcal{F}_t$, the action a^0 will be played repeatedly in T_2 consecutive stages with probability at least $\epsilon_1 := \rho^{n(T_2-1)} > 0$. Supposing this occurs, then at stage $\tau = t + T_2 - 1$, $\text{argmax}_{\alpha_i \in \mathcal{A}_i} \hat{u}_{i,\tau}(\alpha_i) = \text{argmax}_{\alpha_i \in \mathcal{A}_i} u_i(\alpha_i, a^0_{-i}, \mu)$. At this point, either no players can improve their utility (in which case we are at a pure NE), or at least one player can improve their utility. If the latter is the case then, conditioned on $\mathcal{F}_{t+T_2-1}$, with probability at least $\epsilon_2 := \rho^{n-1}(1-\rho)$, exactly one player i chooses to take a best response and improves their utility, and all others continue to play a^0_{-i}. Call the new action profile a^1. Continuing in this manner, we can construct a sequence of actions $a^0, a^1, \ldots, a^m$ (terminating with at most $m = |\mathcal{A}|$) such that a^m is a pure-strategy Nash equilibrium. Conditioned on $\mathcal{F}_t$, the probability of this action sequence occurring (and then the final action a^m being played for T_2 consecutive stages) is bounded from below by $\epsilon := (\epsilon_1\epsilon_2)^{|\mathcal{A}|}\rho^{T_2-1}$. □

The lemma establishes that from any state of beliefs and predictions, the probability that a pure strategy Nash equilibrium is reached and repeated until absorption is positive. Next, we state our main result for this section.

Theorem 1 *Let $\{a_t\}_{t\geq 1}$ be a sequence of actions generated by inertial best response dynamics. Suppose Assumptions 1 - 4 and Condition 1 hold. Then the action sequence $\{a_t\}_{t\geq 1}$ converges to a pure-strategy NE of the game $\Gamma(\mu)$, almost surely. Moreover, let $\tau \in [1, \infty]$ be a random variable indicating the round number in which the action sequence $a(t)$ is absorbed to a pure-strategy NE. Then $E(\tau) < \infty$.*

Proof: Let $t \geq \bar{t}$ and let T_2 be as in Lemma 2. By Lemma 1, if a pure NE action a^* is played in T_2 consecutive stages, then a^* will be played in all consecutive stages. By Lemma 2, the probability of reaching such an "absorbing state" is uniformly lower bounded by some $\epsilon > 0$. Thus, the process is absorbed to a pure NE almost surely in finite time, and $E(\tau) < \infty$ ([20], p.233). □

In the following sections, we will use this result to study the convergence of various instances of inertial best response algorithms in settings with environmental uncertainty and limited communication. In particular our main task will be to show that the algorithms under consideration satisfy Condition 1.

IV. Distributed Implementation of Fictitious Play with Inertia

In this section we study a variation of the classical FP algorithm in which the best response of classical FP is augmented with an inertia term, and inter-agent communication is restricted to a graph. We refer to the class of algorithms studied in this section as Distributed Fictitious Play (D-FP).

The main contributions of this section are twofold. First, we present the general form of the D-FP algorithm (see Section IV-C); this general algorithm explicitly specifies the learning dynamics to be used, but does not explicitly specify the dynamics of the inter-agent communication scheme. We show that any inter-agent communication scheme will lead to pure-strategy NE learning so long as it satisfies a fairly mild condition (see Condition 2). Subsequently, we present an example implementation of D-FP (see Section IV-E) in which the inter-agent communication scheme is explicitly specified and we demonstrate that the example implementation converges to a pure-strategy NE.

We start by reviewing the centralized FP with fading memory and inertia.

A. Fictitious Play with Inertia and Fading Memory

A review of the classical FP algorithm and FP with inertia can be found in [7], [10].

The FP with inertia algorithm is defined as follows. Let $\Psi(a_{i,t}) : \mathcal{A}_i \to \Delta(A_i)$ be the degenerate distribution such that the kth component is $\Psi_k(a_{it}) = 1$ if $a_{it} = k$, and otherwise $\Psi_k(a_{it}) = 0$. Let $f_{i,t} \in \mathcal{R}^{|\mathcal{A}_i|}$ denote the weighted empirical distribution (or just *empirical distribution*) of player i. Formally, let $f_{i,t}$ be defined recursively as $f_{i,1} = \Psi(a_{i,1})$ and for $t \geq 1$,

$$f_{i,t+1} = (1-\alpha)f_{i,t} + \alpha\Psi(a_{i,t+1}), \tag{8}$$

where $\alpha \in (0, 1]$ is a step-size parameter.

Let the joint weighted empirical distribution profile (or joint empirical distribution) be given by $f_t := (f_{1,t}, \ldots, f_{n,t})$. The weighted empirical distribution is said to have "fading memory" because it places greater weight on recent events.[2]

In fictitious play with fading memory and inertia, each player chooses their next-stage action according to the rule

$$a_{i,t+1} \in \begin{cases} \text{argmax}_{\alpha_i \in A_i} u_i(\alpha_i, f_{-i,t}, \theta) & \text{with prob. } 1-\rho, \\ a_{i,t} & \text{with prob. } \rho, \end{cases}$$

where $\rho \in (0, 1)$ is some predefined inertia constant and the probability is conditioned on $\mathcal{F}_{t-1}$ (see Assumption 4).[3]

In the distributed setting, player i may not have precise knowledge of the action history of player j, and hence may not have precise knowledge of $f_{j,t}$. Let $\hat{f}^i_{j,t}$ be an estimate that player i maintains of $f_{j,t}$. Let $\hat{f}^i_t = (\hat{f}^i_{1,t}, \ldots, \hat{f}^i_{n,t})$ be an estimate that player i maintains of the empirical distribution profile f_t.

Next, we present the distributed setup and inter-agent communication scheme.

[2]While the use of inertia is essential to the structure of our proofs, the use of fading memory is less critical. It is possible that the results still hold using a time-varying step size α_t, (e.g., [9], Section II-E); however, the assumption of fading memory simplifies the analysis.

[3]We note that in the classical fictitious play algorithm it is common to omit the state parameter θ from the argument of the utility function, as the implicit assumption is that the players know their own utilities, i.e., $u_i : \mathcal{A} \to \mathbb{R}$. Here we have kept θ to be consistent with our earlier notation.

B. Distributed Setup

In a large-scale setting, the physically distributed nature of the system can make it difficult for players to observe the actions taken by every other player. We assume that players may be unable to observe the actions taken by others, but are equipped with an overlaid communication infrastructure through which they may exchange information with neighboring players. Formally, we assume:

Assumption 5 *Players are equipped with a pre-assigned communication graph $G = (\mathcal{V}, \mathcal{E})$, in which a vertex represents a player, and an edge from vertex i to j represents the ability of player i to communicate information to player j. The directed graph G is strongly connected. Players may exchange information with immediate neighbors (the set of neighbors of a player i is given by $\mathcal{N}_i := \{j \in \mathcal{N} : (i,j) \in \mathcal{E}\}$) once between iterations of the repeated play process.*

C. General D-FP Algorithm

We now state the general D-FP algorithm. The algorithm follows the format of an inertial best response algorithm, but in this case, the utility prediction of player i, $\hat{u}_{i,t}$, is formed using player i's estimate of the empirical distribution $\hat{f}^i_{-i,t}$ (in particular, this is the product distribution of $\{\hat{f}^i_{j,t}\}_{j\neq i}$). Intuitively speaking, this may be thought of as player i forming the (possibly incorrect) belief that her estimate of the empirical distribution $\hat{f}^i_{-i,t}$ accurately represents the mixed strategies of opponents.

Algorithm 2

Initialize

(i) Let $\rho \in (0,1)$ be fixed. For each i, let $a_{i,1}$ be chosen arbitrarily, and let the estimate $\hat{f}^i_1$ be initialized as $\hat{f}^i_1 = \Psi(a_{i,1})$.

Iterate ($t \geq 1$)

(ii) Each agent i chooses their next-stage action according the rule

$$a_{i,t+1} = \begin{cases} \operatorname{argmax}_{\alpha_i \in A_i} u_i(\alpha_i, \hat{f}^i_{-i,t}, \mu^i_t) & \text{with prob. } 1-\rho, \\ a_{i,t} & \text{with prob. } \rho. \end{cases}$$

(iii) Each player i updates their personal empirical distribution $f_{i,t}$ according to (8).

(iv) Each player i engages in one round of information exchange with neighboring agents $j \in \mathcal{N}_i$, and updates their estimate of the joint empirical distribution $\hat{f}^i_t$.

In the following section we analyze convergence of D-FP. We prove convergence to a pure-strategy NE by showing that the algorithm fits the template of Theorem 1.

Remark 1 Note that the framework assumed by Assumptions 1 and 5 may be seen as a strict generalization of the classical repeated play learning framework. The classic framework of centralized and complete information is effectively recovered by letting G be the complete graph and letting $\mu^t_i = \mu = \delta_\theta$ for all $i \in \mathcal{N}$ and $t \geq 1$, where δ_θ is the Dirac delta distribution placing mass 1 on the point $\{\theta\}$.

Remark 2 Note that it has not been precisely specified how $\hat{f}^i_t$ is updated. In the following section we give a general sufficient condition on the update rule under which convergence to pure NE can be guaranteed.

D. Convergence Analysis for Generic D-FP Algorithm

We first introduce some notation necessary for the subsequent development. Given metric space (X, d_X) and a family of functions $\{g_\alpha\}_{\alpha\in A}$ with $g_\alpha : X \to \mathbb{R}$ for all $\alpha \in A$ we say $\{g_\alpha\}_{\alpha\in A}$ is uniformly equicontinuous if for every $\epsilon > 0$ there exists a $\delta > 0$ such that $|g_\alpha(x) - g_\alpha(y)| < \epsilon$, for all $\alpha \in A$ whenever $d_X(x,y) < \delta$, $x, y \in X$.

In order to ensure that weak convergence in beliefs results in a meaningful notion of convergence in terms of players' utility functions, we require that the utility functions satisfy the following assumption:

Assumption 6 *For each $i \in \mathcal{N}$, the family of functions $\{u_i(\cdot,\theta)\}_{\theta\in\Theta}$ is uniformly equicontinuous.*

The following lemma allows us to establish the existence of a finite time such that after this time players' beliefs μ^i_t are sufficiently "close together" for our subsequent analysis.

Lemma 3 *Suppose Assumptions 1 and 6 hold. Let $\epsilon > 0$. There exists a time $T > 0$ such that $|u_i(\sigma, \mu^i_t) - u_i(\sigma,\mu)| < \epsilon$, for all $\sigma \in \Delta^n(\mathcal{A})$, for all i, for all $t \geq T$.*

The proof of the lemma follows immediately from Lemma 9 in the appendix.

Let

$$\eta := \min_{\substack{a'_i, a''_i \in \mathcal{A}_i,\ a'_i \neq a''_i, \\ a_{-i}\in\mathcal{A}_{-i}, \\ i\in\mathcal{N}}} \frac{1}{4}|u_i(a'_i, a_{-i}, \mu) - u_i(a''_i, a_{-i}, \mu)|;$$

the above lemma implies that there exists a time t_η such that, for any $\sigma \in \Delta^n(\mathcal{A})$ and any $i \in \mathcal{N}$ there holds $|u_i(\sigma, \mu^i_t) - u_i(\sigma,\mu)| < \eta$. Intuitively, this means that for all $t \geq t_\eta$, players' beliefs are sufficiently close so that any disparities have negligible effect on the utilities. We note that the existence of such a time t_η is not trivially obtained and relies explicitly on the uniform equicontinuity of the family $\{u_i(\cdot,\theta)\}_{\theta\in\Theta}$. The time t_η here plays the role of time $\bar{t}$ in Section III.

Remark 3 Suppose $t \geq t_\eta$. Then for any player, the best response to any pure strategy is unique; i.e., $\operatorname{argmax}_{\alpha_i\in A_i} u_i(\alpha_i, a_{-i}, \mu^i_t)$ is a singleton for all $a_{-i} \in \mathcal{A}_{-i}$, for all i. This follows immediately from Assumption 3 and the definition of t_η.

The general form of D-FP (see Section IV-C) explicitly specifies the learning dynamics to be used, but does not explicitly specify the dynamics of the inter-agent communication

scheme in step (iii) of the D-FP algorithm. Any protocol for determining $\hat{f}_t^i$ will be valid so long as the following condition is satisfied:

Condition 2 *There exists a $T \in \mathbb{N}_+$ such that if any action $a^* \in \mathcal{A}$ is repeated in $\tilde{T} \geq T$ consecutive stages (i.e., $a_s = a^*,\ s = t, \ldots, t+\tilde{T}-1$), then $\|\hat{f}^i_{t+\tilde{T}-1} - f_{t+\tilde{T}-1}\| < \eta$ for all i.*

We note that the condition is fairly mild and is attainable by a wide variety of distributed tracking algorithms (e.g. [21], [22]). In the following section we present an example implementation of D-FP in which this condition is explicitly satisfied.

The following lemma shows that if Condition 2 is satisfied, then there exists a time T such that if players repeatedly play an action in T consecutive stages, then the set of maximizers of $u_i(\cdot, \hat{f}_t^i, \mu_t^i)$ and $u_i(\cdot, f_t, \mu)$ coincide. The time T here is analogous to the time T in Section III (cf. Condition 1).

Lemma 4 *Let $\{a_t\}_{t\geq 1}$, $\{f_t\}_{t\geq 1}$, and $\{\hat{f}_t^1, \ldots, \hat{f}_t^n\}_{t\geq 1}$ be generated according to Algorithm 2. Suppose Assumptions 1–6, and Condition 2 hold. Assume $t \geq t_\eta$. Then there exists a finite $T \in \mathbb{N}_+$ such that if any action $a^* \in \mathcal{A}$ is repeatedly played in $\tilde{T} \geq T$ consecutive stages, i.e., $a_s = a^*$ for $s = t, \ldots, t+\tilde{T}-1$, then $\operatorname{argmax}_{\alpha_i \in A_i} u_i(\alpha_i, \hat{f}^i_{-i,t+\tilde{T}}, \mu^i_{t+\tilde{T}}) = \operatorname{argmax}_{\alpha_i \in A_i} u_i(\alpha_i, a^*_{-i}, \mu)$ for all i.*

The proof of Lemma 4 is found in Appendix VII.

The preceding lemma ensures that the D-FP process satisfies Condition 1. This allows us to invoke Theorem 1 in order to prove the following theorem, which is the main convergence result for D-FP.

Theorem 2 *Let $\{a_t\}_{t\geq 1}$, $\{f_t\}_{t\geq 1}$, and $\{\hat{f}_t^1, \ldots, \hat{f}_t^n\}_{t\geq 1}$ be generated according to Algorithm 2. Suppose Assumptions 1–6 and Condition 2 hold. Then the action sequence $\{a_t\}_{t\geq 1}$ converges to a pure-strategy Nash equilibrium almost surely.*

Proof: The D-FP process fits the template of Condition 1 with $\hat{u}_{i,t}(\alpha_i) = u_i(\alpha_i, \hat{f}_{-i,t}, \mu_t^i)$ for each i and each $\alpha_i \in \mathcal{A}_i$. By Lemma 4, the sequence of predictions $\{u_i(\alpha_i, \hat{f}_{-i,t}, \mu_t^i)\}_{t\geq 1}$ satisfies Condition 1. Thus, the result follows from Theorem 1. □

In the following section, we present an example implementation of D-FP in which the inter-agent communication scheme is explicitly specified. We demonstrate that the example implementation converges to a pure-strategy NE.

E. An Implementation of D-FP

In this section we provide an example implementation of D-FP in which the mechanism for forming the estimates $\hat{f}_t^i$ and the distributed information dissemination scheme are explicitly specified.

Algorithm 3
Initialize

(i) Let $\rho \in (0,1)$ be fixed. For each i, let $a_{i,1}$ be chosen arbitrarily. Let $f_{i,1} = \Psi(a_{i,1})$ for all i and let $\hat{f}^i_{i,1} = f_{i,1}$. For $j \neq i$, let $\hat{f}^i_{j,1} = w^i_{j,j} f_{j,1}$ if $j \in \mathcal{N}_i$ and $\hat{f}^i_{j,1} = 0$ otherwise, where $w^i_{j,j}$ is a weight constant (see step (iv) and Lemma 5).

Iterate ($t \geq 1$)
(ii) Each agent i chooses their next-stage action according the rule

$$a_{i,t+1} = \begin{cases} \operatorname{argmax}_{\alpha_i \in A_i} u_i(\alpha_i, \hat{f}^i_{-i,t}, \mu^i_t) & \text{with prob. } 1-\rho, \\ a_{i,t} & \text{with prob. } \rho. \end{cases}$$

(iii) Each player i updates their personal empirical distribution $f_{i,t}$ according to (8).

(iv) For each i, $\hat{f}^i_{j,t}$ is updated as

$$\hat{f}^i_{j,t+1} = \sum_{k \in \mathcal{N}_i} w^i_{j,k} \left(\hat{f}^k_{j,t} + (f_{j,t+1} - f_{j,t})\, \chi_{\{k=j\}} \right), \quad (9)$$

where $\chi_{\{k=j\}}$ is the characteristic function defined by $\chi_{\{k=j\}} = 1$ if $k = j$ and $\chi_{\{k=j\}} = 0$ otherwise, and where $w^i_{j,k}$ is the weight that player i attributes to k's estimate of j's empirical frequency (see Lemma 5.)

F. An Implementation of D-FP: Convergence Analysis

In this section we analyze the convergence of the example implementation of D-FP presented in Section IV-E.

The following Lemma ensures that Condition 2 is satisfied. It then follows by Theorem 2 that the example implementation of D-FP converges to a pure-strategy NE.

Lemma 5 *Let $\{a_t\}_{t\geq 1}$, $\{f_t\}_{t\geq 1}$, and $\{\hat{f}_t^1, \ldots, \hat{f}_t^n\}_{t\geq 1}$ be generated according to Algorithm 3. Suppose Assumption 5 holds. Let $W_j \in R^{n\times n}$, $j \in \mathcal{N}$ be a* weight matrix *with the i,k-th entry given by $W_j(i,k) = w^i_{j,k}$. Assume that the matrix W_j is row stochastic with sparsity conforming to the communication network G. Assume the j-th diagonal entry satisfies $w^j_{j,j} = 1$ for each $j \in \mathcal{N}$. Let P_j be the matrix obtained by removing the j-th row and column from W_j. Assume P_j is irreducible and substochastic in the sense that at least one row sum of P_j is strictly less than* 1.

Then, for any $\epsilon > 0$ there exists $T \in \mathbb{N}_+$ such that if players repeat any action $a^* \in \mathcal{A}$ for $\tilde{T} \geq T$ consecutive stages (i.e., $a_s = a^*,\ s = t, \ldots, t+\tilde{T}-1$) then $\|\hat{f}^i_{t+\tilde{T}-1} - f_{t+\tilde{T}-1}\| < \epsilon$.

The proof of Lemma 5 is found in the appendix.

The size of the T term in the above lemma depends on the step size parameter α in (8), network structure G, and the weight matrices $\{W_j\}_{j\in\mathcal{N}}$. In general, we expect T to be small if the step size is close to 1, $0 \ll \alpha \leq 1$, and the diameter of the network is small. For instance, in a complete network with $\alpha = 1$ and weights $w^i_{k,k} = 1$ if $k \in \mathcal{N}_i$ we have $T = 1$ time steps. In a star network with $\alpha = 1$ such that the weights associated with the center player i satisfy $w^i_{k,k} = 1$ if $k \in \mathcal{N}_i$ and the weights associated with fringe players satisfy $j \neq i$ $w^j_{k,i} = 1$ for $k \neq j$, then we have $T = 2$ time steps.

In the D-FP implementation presented here, each player i maintains an estimate of the empirical distribution of *every* individual player, and exchanges these estimates with its neighbors $\{\hat{f}^i_{j,t}\}_{j\in\mathcal{N}}$. This means that at each time step players exchange $\sum_{j=1}^n |\mathcal{A}_j|$ values, and in general, track $\sum_{j=1}^n |\mathcal{A}_j|$ values.

In the following section we study a distributed inertial best response algorithm that can alleviate this communication and memory burden within a certain class of games.

V. Distributed Implementation of Joint Strategy Fictitious Play with Inertia

Joint Strategy FP (JSFP) with inertia—introduced in [9]—is a variant of FP developed for large-scale games that has relatively low computational complexity and low information overhead requirements. In this section we study a distributed implementation of JSFP with inertia (referred to hereafter as D-JSFP) for use in settings with environmental uncertainty. The algorithms in this section may be seen as instances of inertial best response dynamics.[4]

The variant of JSFP that we study is applicable within the class of congestion games—a subset of the more general class of weakly-acyclic games (see Assumption 2). This restriction comes as a consequence of the aggregative information that is shared over the communication graph. Thus, while D-JSFP operates with lower complexity and communication overhead than D-FP (Section IV), D-JSFP is applicable within a narrower class of games than D-FP.

We begin by introducing the class of congestion games in Section V-A. We then introduce the general D-JSFP algorithm in Sections V-B–V-C. The general D-JSFP algorithm may be implemented using a wide range of inter-agent communication schemes. We provide sufficient conditions for the convergence of the general D-JSFP algorithm in Secion V-D, and prove convergence in the same section. Finally, in Sections V-E–V-F we provide an explicit example implementation of D-JSFP in which the inter-agent communication dynamics are explicitly defined, and prove convergence of the algorithm.

A. Congestion Games

Let $R = \{1, \dots, m\}$ denote a set of resources. For each $i \in \mathcal{N}$, let $\mathcal{A}_i \subseteq 2^R$, where 2^R denotes the power set of R. In particular, an action choice a_i indicates a subset of resources being utilized by player i.

In a congestion game, the cost associated with using a resource is dependent on the total number of players using the same resource. For each $r \in R$, $a \in \mathcal{A}$, let $N_r(a) \in \mathbb{N}$ denote the number of players using resource r under the action profile a. More generally, for a subset of players $\mathcal{K} \subseteq \mathcal{N}$, the number of players in $\mathcal{K}$ utilizing resource r given $\{a_j\}_{j\in\mathcal{K}}$, is given by

$$N_r(\{a_j\}_{j\in\mathcal{K}}) := \sum_{j\in\mathcal{K}} \mathbf{1}(r \in a_j).$$

[4]A related variant of JSFP—termed Average Strategy FP (ASFP) is studied in [23]. However, ASFP differs fundamentally from D-JSFP in that (i) ASFP assumes instantaneous and perfect information dissemination by an oracle, (ii) D-JSFP uses a projection operation to make sense of the notion of players "assuming" that the average congestion profile represents choices taken by agents, and (iii) D-JSFP considers learning under environmental uncertainty.

where $\mathbf{1}(r \in a_j) = 1$ if $r \in a_j$ and $\mathbf{1}(r \in a_j) = 0$ otherwise. Given a subset of players $\mathcal{K}$, and a corresponding set of actions $\{a_j\}_{j\in\mathcal{K}}$, we represent the number of players using each resource by $N(\{a_j\}_{j\in\mathcal{K}})$, where $N : \prod_{j\in\mathcal{K}} A_j \to \mathbb{N}^m$ is a mapping with the r-th entry in $N(\{a_j\}_{j\in\mathcal{K}})$ given by $N_r(\{a_j\}_{j\in\mathcal{K}})$.

For $r \in R$ and $k \in \mathbb{N}$, let $c_r(k, \theta_r)$ be the cost associated with using resource r given state θ_r, when there are precisely k players simultaneously using the resource. For $a_i \in \mathcal{A}_i$ and $N_r(a_{-i}) \in \mathbb{N}$ and state parameter $\theta := \{\theta_r\}_{r\in R} \in \Theta$, let the utility of player i be given by

$$u_i(a_i, a_{-i}, \theta) = -\sum_{r\in a_i} c_r(N_r(a_{-i}) + N_r(a_i), \theta_r)$$

where we have written $N_r(a) = N_r(a_{-i}) + N_r(a_i)$ explicitly to emphasize dependence of the utility on "self action" a_i and actions of other players a_{-i}. Note that within the class of congestion games, players do not need to precisely know the full action profile $a = (a_1, \dots, a_n) \in \mathcal{A}$ to compute their utility. It is sufficient for each player to have knowledge of $N(a_{-i}) \in \mathbb{N}^m$ and their own action $a_i \in \mathcal{A}_i$. In this context, we sometimes express the utility function using the abuse of notation $u_i(a_i, N(a_{-i}), \theta) = u_i(a_i, a_{-i}, \theta)$.

In the following, we use this property of the utility functions in congestion games to design the D-JSFP algorithm which has a lower communication overhead than D-FP.

B. D-JSFP Setup

Assume players repeatedly face off in a congestion game. We define $\zeta_{i,t}(r)$ to be a (fading-memory) weighted average used to track the amount of congestion induced on resource r by the actions of (only) player i. In particular, let $\zeta_{i,t}(r)$ be defined recursively by $\zeta_{i,1}(r) := N_r(a_{i,1})$, and for $t \geq 1$,

$$\zeta_{i,t+1}(r) := (1-\alpha)\zeta_{i,t}(r) + \alpha N_r(a_i(t)), \tag{10}$$

where $\alpha \in (0, 1]$ is a weight parameter inducing a fading-memory effect (cf. (8) and subsequent discussion).

Furthermore, define $\zeta_{i,t} \in \mathbb{R}^m$ to be the vector stacking $(\zeta_{i,t}(r))_{r\in R}$—this is a vectorized representation of the congestion induced by player i on any given resource.

Define $\zeta_t(r) := \sum_{j\in\mathcal{N}} \zeta_{j,t}(r)$—this represents the congestion induced on resource r by the actions of *all* players. Note this can also be expressed recursively as $\zeta_t(r) = (1-\alpha)\zeta_t(r) + \alpha N_r(a(t))$.

Similar to the above, let ζ_t be a vector in $\mathbb{R}^m$ stacking $(\zeta_t(r))_{r\in R}$—this is a vectorized representation of the congestion induced by *all* players on any given resource. We refer to ζ_t as the empirical congestion distribution.

In the distributed framework, players may not have precise knowledge of ζ_t. Instead, we assume each player i maintains an estimate of ζ_t which we denote by $\hat{\zeta}^i_t \in \mathbb{R}^m$. The r-th term of player i's estimate, $\hat{\zeta}^i_t(r)$, represents her estimate of the congestion at resource $r \in R$.

Finally, in order to rigorously define D-JSFP, we require the following notion of a projection. For a vector $v \in \mathbb{R}^m$ define $P(v)$ to be a projection of v onto the set of non-negative m-dimensional integer-valued vectors $\mathbb{N}^m$; formally, for $1 \leq$

$r \leq m$, let $P(v)(r) := z$ for the unique $z \in \mathbb{N}$ satisfying $z - \frac{1}{2} \leq v(r) < z + \frac{1}{2}$. Let $P(v)$ be the vector stacking $\{P(v)(r)\}_{r \in R}$.

C. General D-JSFP Algorithm

We now state the general D-JSFP algorithm. The algorithm follows the format of an inertial best response algorithm, but in this case, the utility prediction of player i $\hat{u}_{i,t}$ is formed using player i's estimate of the empirical congestion distribution $\hat{\zeta}_t^i$. The algorithm is formally stated below.

Algorithm 4

initialize

(i) Let $\rho \in (0,1)$ be the inertia probability. Let each player i choose an arbitrary initial action $a_{i,1}$. Initialize the estimate $\hat{\zeta}_1^i$ and belief on state μ_1^i.

iterate $(t \geq 1)$

(ii) Let $\hat{\zeta}_{-i,t}^i := \hat{\zeta}_t^i - \zeta_{i,t}$. For each player i, the next-stage action is chosen according to the rule

$$a_{i,t+1} \in \begin{cases} \operatorname{argmax}_{\alpha_i \in \mathcal{A}_i} u_i(\alpha_i, P(\hat{\zeta}_{-i,t}^i), \mu_t^i), & \text{w.p. } 1-\rho \\ a_{i,t}, & \text{w.p. } \rho \end{cases}$$

(iii) Each player i updates their state belief μ_t^i, exchanges information with neighboring agents $j \in \mathcal{N}_i$ and updates their estimate $\hat{\zeta}_t^i$.

Remark 4 The D-JSFP algorithm is closely related to the JSFP algorithm proposed in [9]. In particular, the vector representing the weighted congestion at each resource $\zeta_t := (\zeta_t(r))_{r \in \mathcal{N}}$ is closely related to the empirical frequency of the joint action profile used in the JSFP algorithm. In this case, the vectors ζ_t and $\zeta_{i,t}$ together may be thought of as a sufficient statistic to compute the predicted utility. We note that in D-JSFP, the complexity problem associated with FP in large games is mitigated in a different manner than classical JSFP. In classical JSFP, the expected utility is directly tracked using a simple recursion. In our formulation of JSFP, the utility is not directly tracked; instead, players track some "sufficient statistic" $\zeta_{-i,t} = \zeta_t - \zeta_{i,t}$ and only evaluate the utility at nearby pure strategies via the projection operation, $u_i(a_i, P(\zeta_t - \zeta_{i,t}), \mu_t^i)$.

D. Convergence Analysis for General D-JSFP Algorithm

The general form described in Section V-C explicitly specifies the learning dynamics to be used, but does not explicitly specify the dynamics of the inter-agent communication scheme. In the following we show that any inter-agent communication scheme will lead to convergence so long as it satisfies the following condition (cf. Condition 2 for D-FP).

Condition 3 *There exists a $T \geq 1$ such that if any action a^* is repeated in $\tilde{T} \geq T$ consecutive stages then $|\hat{\zeta}_{t+\tilde{T}}^i(r) - N_r(a^*)| < \frac{1}{4}$ for every $r \in R$.*

Given Condition 3, we show in the following lemma that if any action a^* is repeated in sufficiently many stages, then each player's estimate $\hat{\zeta}_{-i,t}^i$ may be brought sufficiently close to the congestion profile $N(a_{-i}^*)$ to ensure convergence.

Lemma 6 *Let $\{a_t\}_{t\geq 1}$ and $\{\hat{\zeta}_t^1, \ldots, \hat{\zeta}_t^n\}_{t \geq 1}$ be generated according to a D-JSFP process as defined in Section V-C, and let $\{\zeta_{i,t}\}_{i \in \mathcal{N}, t \geq 1}$ be as defined in* (10)*. Assume Condition 3 holds. There exists a $T \geq 1$ such that if any action a^* is repeated in $\tilde{T} \geq T$ consecutive stages then $|\hat{\zeta}_{-i,t+\tilde{T}}^i(r) - N_r(a_{-i}^*)| < \frac{1}{2}$ for every $r \in R$.*

Proof: Let $\tilde{T} \geq 0$ and note that

$$\begin{aligned} &|\hat{\zeta}_{-i,t+\tilde{T}}^i(r) - N_r(a_{-i}^*)| \\ &= |(\hat{\zeta}_{t+\tilde{T}}^i(r) - \zeta_{i,t+\tilde{T}}(r)) - (N_r(a^*) - N_r(a_i^*))| \\ &\leq |\hat{\zeta}_{t+\tilde{T}}^i(r) - N_r(a^*)| + |\zeta_{i,t+\tilde{T}}(r) - N_r(a_i^*)|. \end{aligned} \tag{11}$$

By Condition 3, we may choose T' such that if a^* is repeated in $\tilde{T} \geq T'$ consecutive stages, there holds $|\hat{\zeta}_{t+\tilde{T}}^i(r) - N_r(a^*)| < \frac{1}{4}$. Note also that $|\zeta_{i,t+\tilde{T}}(r) - N_r(a_i^*)| \to 0$ as $\tilde{T} \to \infty$ (this follows from (10)), and thus we may choose T'' such that for $\tilde{T} \geq T''$ there holds $|\zeta_{i,t+\tilde{T}}(r) - N_r(a_i^*)| < \frac{1}{4}$. Letting $T = \max\{T', T''\}$, the desired result follows from (11). □

The next lemma sets us up to prove convergence of D-JSFP using Theorem 1 and Condition 1.

Lemma 7 *Let $\{a_t\}_{t\geq 1}$ and $\{\hat{\zeta}_t^1, \ldots, \hat{\zeta}_t^n\}_{t \geq 1}$ be generated according to a D-JSFP process as defined in Section V-C, and let $\{\zeta_{i,t}\}_{i \in \mathcal{N}, t \geq 1}$ be as defined in* (10)*. Suppose Assumptions 1–5 and Condition 3 hold. There exists a $T \geq 1$ such that if any action a^* is repeated in $\tilde{T} \geq T$ consecutive stages then $\operatorname{argmax}_{\alpha_i \in \mathcal{A}_i} u_i(\alpha_i, P(\hat{\zeta}_{-i,t+\tilde{T}}^i), \mu_{t+\tilde{T}}^i) = \operatorname{argmax}_{\alpha_i \in \mathcal{A}_i} u_i(\alpha_i, N(a_{-i}^*), \mu)$.*

Proof : Let T' be chosen as in Lemma 6 so that $|\hat{\zeta}_{-i,t+\tilde{T}}^i(r) - N_r(a_{-i}^*)| < \frac{1}{2}$ for every $r \in R$, $i \in \mathcal{N}$ and all $\tilde{T} \geq T'$. It follows that $P(\hat{\zeta}_{-i,t+\tilde{T}}^i) = N(a_{-i}^*)$. Thus, $\operatorname{argmax}_{\alpha_i \in \mathcal{A}_i} u_i(\alpha_i, P(\hat{\zeta}_{-i,t+\tilde{T}}^i), \mu_{t+\tilde{T}}^i) = \operatorname{argmax}_{\alpha_i \in \mathcal{A}_i} u_i(\alpha_i, N(a_{-i}^*), \mu_{t+\tilde{T}}^i)$. Moreover, there exists a $\bar{t}$ such that for $t \geq \bar{t}$ there holds $\operatorname{argmax}_{\alpha_i \in \mathcal{A}_i} u_i(\alpha_i, a_{-i}, \mu_{t+\tilde{T}}^i) = \operatorname{argmax}_{\alpha_i \in \mathcal{A}_i} u_i(\alpha_i, a_{-i}, \mu)$ for all $a_{-i} \in \mathcal{A}_{-i}$. This is a consequence of Assumption 1 and the fact that u_i is continuous and bounded, and that $\mathcal{A}_{-i}$ is a finite set. By setting $T = \max\{\bar{t}, T'\}$ we get the desired result. □

In the following theorem we combine the results above with Theorem 1 to prove that the D-JSFP algorithm converges to a pure-strategy Nash equilibrium.

Theorem 3 *Let $\{a_t\}_{t\geq 1}$ and $\{\hat{\zeta}_t^1, \ldots, \zeta_t^n\}_{t \geq 1}$ be generated according to a D-JSFP process as defined in Section IV-C, and let $\{\zeta_t\}_{t \geq 1}$ be as defined in* (10)*. Suppose Assumptions 1–5 and Condition 3 hold. Then the action sequence $\{a_t\}_{t\geq 1}$ converges to a pure strategy Nash equilibrium, almost surely.*

Proof : The process fits the template of Theorem 1, with $u_i(\alpha_i, P(\hat{\zeta}_{-i,t}^i), \mu_t^i) = \hat{u}_{i,t}(\alpha_i)$ for each i and each $\alpha_i \in \mathcal{A}_i$.

By Lemma 7, the sequence $\{u_i(\alpha_i, P(\zeta^i_{-i,t}), \mu^i_t)\}_{t\geq 1}$ satisfies Condition 1. The result then follows from Theorem 1. □

E. An Implementation of D-JSFP

In this section we provide an example implementation of D-JSFP in which the mechanism for forming the estimate $\hat{\zeta}^i_t$, and the distributed information dissemination scheme are explicitly specified. The algorithm is as follows:

Algorithm 5
initialize
(i) Let $a_{i,1}$ be arbitrary for all i. Let $\hat{\zeta}^i_1 = N(a_{i,1})$ for all i.
iterate $(t \geq 1)$
(ii) Let $\hat{\zeta}^i_{-i,t} = \hat{\zeta}^i_t - \zeta_{i,t}$. For each player i, the next-stage action is chosen according to the rule

$$a_{i,t+1} \in \begin{cases} \text{argmax}_{\alpha_i \in \mathcal{A}_i} u_i(\alpha_i, P(\hat{\zeta}^i_{-i,t}), \mu^i_t), & \text{w.p. } 1-\rho \\ a_{i,t}, & \text{w.p. } \rho. \end{cases}$$

(iii) Update $\zeta_{i,t+1}$ according to (10).
(iv) Each player i updates their estimate of ζ_t as:

$$\hat{\zeta}^i_{t+1} := \sum_{k \in \mathcal{N}_i} w^i_k \left(\hat{\zeta}^k_t + \zeta_{k,t+1} - \zeta_{k,t} \right)$$

where w^i_k denotes the weight that player i places on the information received from player k (see Lemma 8).

We remark that in D-JSFP players only share a vector with m integer values with their neighbors where m is the cardinality of the set of resources R. In comparison, in the D-FP algorithm players share their estimate of each agent's empirical frequency implying that they share $n \times m$ values at each step.

F. Convergence Analysis of Example Implementation of D-JSFP

The following lemma shows that the distributed information tracking scheme used to form the estimates $\hat{\zeta}^i_t$ (see step (iii) in Section V-E) satisfies Condition 3 given an appropriate assumption on players' weights w^i_k. It then follows by Theorem 3 that the given implementation of D-JSFP converges to a pure strategy Nash equilibrium.

Lemma 8 *Let $\{a_t\}_{t\geq 1}$ and $\{\hat{\zeta}^1_t, \ldots, \zeta^n_t\}_{t\geq 1}$ be generated according to a D-JSFP process as defined in Section V-E, and let $\{\zeta_t\}_{t\geq 1}$ be as defined in* (10)*. Suppose Assumptions 1–5 hold. Let $W \in R^{n\times n}$ be a* weight matrix *with the i,kth entry given by $W(i,k) = w^i_k$. Assume that the matrix W is doubly stochastic, aperiodic, and irreducible.*

There exists a $\tilde{T} \geq 1$ such that if any action a^* is repeated in $\tilde{T} \geq T$ consecutive stages then $|\hat{\zeta}^i_{t+\tilde{T}}(r) - N_r(a^*)| < \frac{1}{4}$ for every $r \in R$.

The proof of Lemma 8 is similar to the proof of Lemma 5, and is omitted here for brevity.

VI. Numerical examples

A. Traffic congestion game

We consider the congestion game defined in Section V-A. We assume each agent can only pick a single route at each time $\mathcal{A}_i = R$. The cost of route $r \in R$ is quadratic in the number of players on that route,

$$c_r(\sigma_r(a)) = \phi_r \sigma_r(a)^2 + \beta_r \sigma_r(a) + \kappa_r$$

where $\phi_r, \beta_r, \kappa_r$ are route specific constants which play the role of the unknown state of the world.

In the numerical setup, we let $\phi_r = 1/(1+2(r-1))$, $\beta_r = 1$ and $\kappa_r = 1$ for each route. Agents implement the D-JSFP algorithm defined in Section V-E. We let the fading constant α and the inertia probability ρ be 0.1 and 0.3, respectively. Players are uncertain about κ_r and ϕ_r. Each player receives a single initial signal on κ_r and ϕ_r for each $r \in R$. The signals on κ_r come from the uniform distribution in $[0.75, 1.25]$. The signals on ϕ_r come from the uniform distribution in $[\phi_r - 0.25, \phi_r + 0.25]$. Following each round of game play, players exchange their mean estimates of the state parameters κ_r and ϕ_r with neighbors, and form new estimates using a weighted average consensus recursion.

In this simple example, the μ^i_t are deterministic delta distributions, as is the limit distribution μ. As noted in Section III, this achieves weak convergence of the beliefs, thus satisfying Assumption 1. Moreover, Assumption 1 could also be satisfied in an almost sure sense by analogously generating a sequence of *random* delta distributions using consensus recursions that ensure almost sure convergence of the estimates [24]–[26].

There are $n = 10$ agents and $m = 3$ routes. The connection among agents is determined by a geometric network. Agents are randomly placed on a two dimensional space (5 units $\times$ 5 units) and two agents are connected if their distance is less than a given threshold value (2 units). In Fig. 1 we plot the number of vehicles on each route. In Fig. 2 we plot the congestion cost at each road and total congestion cost over time. We initialize all agents' actions to be route 1. We observe that at equilibrium there are four agents on roads 2 and 3, and two agents on road 1. Total congestion cost diminishes significantly in the first few time steps – see Fig. 2. At equilibrium we observe that congestion cost at road 2 is slightly higher than the other roads.

Next, we analyze convergence times with respect to initial level of uncertainty on the state parameters κ_r and ϕ_r, and network connectivity. In particular we use distance thresholds 2, 2.5, 3 and 3.5 to generate four geometric networks with diameters 4, 3, 2 and 2 and with average lengths 2, 1.3, 1.2 and 1.1, respectively. For each network, we make 50 runs with and without uncertainty on the state parameters. The average convergence time of the D-JSFP algorithm is reported in Table I. We observe convergence time tends to decrease as the connection threshold increases. We do not see a significant difference between the convergence times when there is uncertainty or not about the state parameters.

B. Target assignment game

We consider a team of n robots, and n target objects. Each robot can target one object and the team's goal is to target all of

	connection threshold (units)			
	2	2.5	3	3.5
Certain	361	172	220	104
Uncertain	329	179	212	95

TABLE I
AVERAGE CONVERGENCE TIME OF THE D-JSFP ALGORITHM

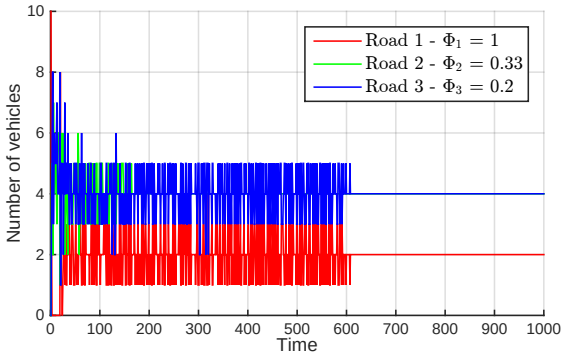

Fig. 1. Number of vehicles on each route. There are $n = 10$ vehicles each day choosing among $m = 3$ possible routes. Convergence to pure Nash equilibrium happens after time 600.

the objects. The action space is the set of objects $\{1, 2 \ldots, n\}$ for each robot. The payoff of a robot i targeting object k is inversely proportional to its distance to the object, represented by $d(i, k)$, if no other robot is targeting object k,

$$u_i(a_i = k, a_{-i}, d(i,k)) = d(i,k)^{-1}\mathbf{1}(\sum_{j=1}^{n}\mathbf{1}(a_j = k) = 1)$$

where $\mathbf{1}(\cdot)$ is the indicator function. The target assignment game with payoffs as above is a congestion game with each object representing a resource. Note that any action profile that covers all the objects is a Nash equilibrium because any unilateral deviation from such profile results in zero payoff for the deviating agent. The optimal Nash equilibrium profile minimizes the total distance while covering all the objects.

We assume robots do not exactly know their distance to each target. Robots receive an initial noisy signal about the location

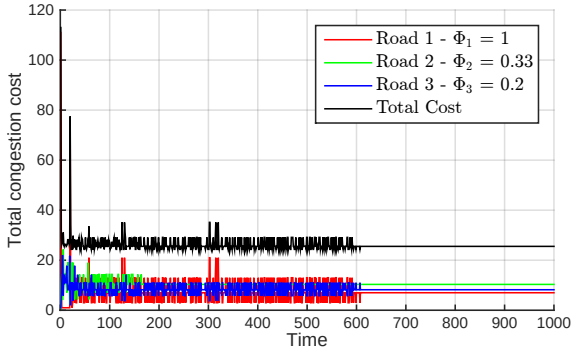

Fig. 2. Congestion cost at each road and total congestion cost $\sum_r c_r$ at each time. Initial total cost dampens quickly in the first few time steps.

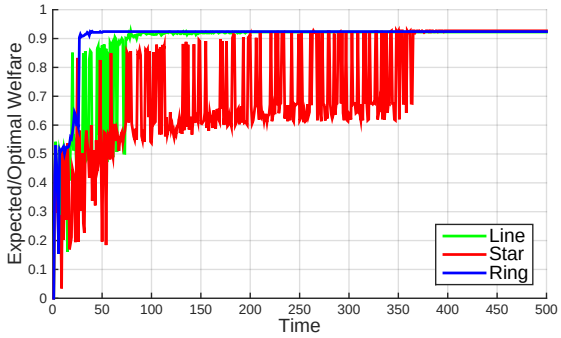

Fig. 3. Expected welfare normalized by optimal welfare for line, star and ring networks. The expected welfare is the same for all communication networks. Convergence time to a Nash equilibrium is the fastest for the ring network and slowest for the star network.

of each object $\tilde{\theta}_i(k)$ distributed according to $\mathcal{N}(\theta(k), \xi_k)$ where $\theta(k)$ is the location of the object and $\xi_k > 0$ is the variance of the noisy signal. Robot i updates its belief about each object by averaging its neighbors' beliefs with initial belief equal to $\tilde{\theta}_i(k)$. The distance to an object is the norm of the difference between the locations of the robot and the object.

In the numerical setup, we consider $n = 5$ robots and $n = 5$ objects. For comparison, we consider line, ring and star communication networks. For each network, we consider 50 runs with random initialization of the signals. In Figure 3, we plot sample average welfare normalized by the optimal welfare over time for each network type. Welfare at time t is defined as the sum of utilities of robots given the action profile generated by the D-JSFP process at time t. The expected welfare at time t is the average of welfare values at time t obtained over 50 runs. Optimal welfare is the value of welfare obtained by the action profile that maximizes the welfare. Fig. 3 shows that the expected welfare of the Nash equilibrium reached by the algorithm is the same regardless of the communication network. However, the convergence time of the algorithm depends on the network structure where the star network has the slowest convergence time and the ring network has the fastest convergence time.

VII. CONCLUSION

We have studied general inertial best response dynamics for learning pure-strategy NE in distributed multi-agent systems with uncertainty. Subsequently, we have studied two important cases of inertial best response dynamics (FP and JSFP); we have derived distributed variants of both algorithms, derived sufficient conditions for convergence, and provided explicit example implementations. Finally, results were corroborated with simulation examples in congestion games.

APPENDIX

The following lemma studies a uniform convergence property for families of equicontiuous functions. The lemma is applied in the text to prove that given weak convergence of beliefs (see Assumption 1) and equicontinuity of utilities (see Assumption 6) Lemma 3 holds.

Lemma 9 *Let Θ be as defined in Section II. Let X be a compact subset of a finite-dimensional Euclidean space. Let $g : X \times \Theta \to \mathbb{R}$ be bounded. Assume that for each $x \in X$ the function $g(x,\cdot) : \Theta \to \mathbb{R}$ is Borel measurable. Furthermore, assume that the family of functions $\{g(\cdot,\theta)\}_{\theta\in\Theta}$ is uniformly equicontinuous. Let $(\mu_t)_{t\geq 1}$ and μ be probability measures on $(\Theta, \mathcal{B}(\Theta))$ such that $\mu_t \xrightarrow{w} \mu$. For any $\epsilon > 0$ there exists $T \in \mathbb{N}_+$ such that for all $t \geq T$ there holds*

$$|\int_\Theta g(x,\theta)\mu_t(d\theta) - \int_\Theta g(x,\theta)\mu(d\theta)| < \epsilon, \text{ for all } x \in X.$$

Proof: Let $\epsilon > 0$. Since the family of functions $\{g(\cdot,\theta)\}_{\theta\in\Theta}$ is equicontinuous, there exists a $\delta > 0$ such that if $x, y \in X$ and $\|x - y\| < \delta$ then $|g(x,\theta) - g(y,\theta)| < \frac{\epsilon}{3}$ for any $\theta \in \Theta$. Since X is a compact subset of Euclidean space we may choose a finite set $\{x_n\}_{n=1}^N$ such that $\bigcup_{n=1}^N B_\delta(x_n) \supset X$, where $B_\delta(x_0)$ denotes the open ball of radius δ centered at x_0.

Since μ_t converges weakly to μ, and the set $\{x_n\}_{n=1}^N$ is finite, and the function $g(x',\cdot)$ is measurable for each $x' \in \{x_n\}_{n=1}^N$, there exists a $T \in \mathbb{N}_+$ such that

$$|\int g(x',\theta)\mu_t(d\theta) - \int g(x',\theta)\mu(d\theta)| < \frac{\epsilon}{3}$$

for all $x' \in \{x_n\}_{n=1}^N$.

Choose an arbitrary $x \in X$. There exists a $\tilde{x} \in \{x_n\}_{n=1}^N$ such that $x \in B_\delta(\tilde{x})$. In particular, note that

$$|\int [g(x,\theta) - g(\tilde{x},\theta)]\mu_t(d\theta)| < \frac{\epsilon}{3}\int \mu_t(d\theta) = \frac{\epsilon}{3},$$

and likewise,

$$|\int [g(x,\theta) - g(\tilde{x},\theta)]\mu(d\theta)| < \frac{\epsilon}{3}\int \mu(d\theta) = \frac{\epsilon}{3}.$$

Using the triangle inequality and applying the estimates derived above gives,

$$\begin{aligned}
&|\int g(x,\theta)\mu_t(d\theta) - g(x,\theta)\mu(d\theta)| \\
&\leq |\int [g(x,\theta) - g(\tilde{x},\theta)]\mu_t(d\theta)| \\
&+ |\int g(\tilde{x},\theta)\mu_t(d\theta) - \int g(\tilde{x},\theta)\mu(d\theta)| \\
&+ |\int [g(\tilde{x},\theta) - g(x,\theta)]\mu(d\theta)| < \epsilon,
\end{aligned}$$

for all $t \geq T$. □

The following is a proof of Lemma 4 in Section IV-D.

Proof (Lemma 4): We first show that for every $\xi > 0$, there exists a finite $T' \in \mathbb{N}_+$ such that if any action $a^* \in \mathcal{A}$ is repeatedly played in $T \geq T'$ consecutive stages, i.e., $a_s = a^*$ for $s = t, \ldots, t+T$, then $\|\hat{f}^i_{t+T} - \Psi(a^*)\| < \xi$, for all i.

Let $\xi > 0$ be given. By the triangle inequality, $\|\hat{f}^i_t - \Psi(a^*)\| \leq \|\hat{f}^i_t - f_t\| + \|f_t - \Psi(a^*)\|$. Thus it is sufficient to choose T' sufficiently large so that

$$\|\hat{f}^i_t - f_t\| < \frac{\xi}{2} \text{ and } \|f_t - \Psi(a^*)\| < \frac{\xi}{2}, \text{ for all } t \geq T'. \quad (12)$$

Claim there exists a finite $T'_1 > 0$ such that if any action a^* is repeated in T'_1 consecutive stages, then $\|f_{t+T'_1} - \Psi(a^*)\| < \frac{\xi}{2}$. This follows immediately from the observation that, for all i, and for any $\tau \in \{1, 2, \ldots\}$ there holds $f_{i,t+\tau} = \alpha(1-(1-\alpha)^{\tau-1})\Psi(a^*_i) + (1-\alpha)^\tau f_{i,t}$.

By Condition 2 there exists a T'_2 such that if any action a^* is repeated for $T \geq T'_2$ consecutive stages, then $\|\hat{f}^i_{t+T} - f_{t+T}\| < \frac{\xi}{2}$. Let $T' = \max(T'_1, T'_2)$ and note that for $t \geq T'$, (12) holds.

We now move to prove the statement in the Theorem. Note that under Condition 2, there exists a compact C such that for each player i the belief $\hat{f}^i_t$ belongs to C for all t. For each i, $u_i(\cdot,\mu)$ (as defined in (3)) is multilinear in the first argument. Hence, it is locally Lipschitz continuous in the first argument; i.e., given the belief μ there exists a finite $K_i > 0$ such that for any $p', p'' \in C$ there holds $|u_i(p',\mu) - u_i(p'',\mu)| \leq K_i\|p' - p''\|$. Let $K := \max\{K_i : i \in \mathcal{N}\}$.

Since the action space is finite, Assumption 3 implies there exists an $\epsilon > 0$ such that for all $i \in \mathcal{N}$ there holds $|u_i(a'_i, a_{-i},\mu) - u_i(a''_i, a_{-i},\mu)| > \epsilon$ for all $a'_i, a''_i \in \mathcal{A}_i$, $a'_i \neq a''_i$ $a_{-i} \in \mathcal{A}_{-i}$.

Let $\xi < \frac{\epsilon}{4K}$. Applying the claim demonstrated in the beginning of the proof there exists a $T > 0$ such that if any action $a^* \in \mathcal{A}$ is repeated in $\tilde{T} \geq T$ consecutive stages, then $\|\hat{f}^i_{t+T} - \Psi(a^*)\| < \xi$, for all i.

Let $\{a'_i\} = \operatorname{argmax}_{\alpha_i\in A_i} u_i(\alpha_i, a^*_{-i},\mu)$ (by Remark 3, the best response to any pure strategy is a singleton), and note that

$$u_i(a'_i, a^*_{-i},\mu) - \epsilon > u_i(a''_i, a^*_{-i},\mu) \quad (13)$$

for all $a''_i \in \mathcal{A}_i, a''_i \neq a'_i$. Furthermore, for $\tau = t + \tilde{T}$, $\tilde{T} \geq T$, there holds by Lipschitz continuity of $u_i(\cdot,\mu)$,

$$|u_i(a_i, \hat{f}^i_{-i,\tau},\mu) - u_i(a_i, a^*_{-i},\mu)| \leq K\|\hat{f}^i_\tau - a^*\| < \frac{\epsilon}{4},$$

for all $a_i \in \mathcal{A}_i$. In particular, for $\{a'_i\} = \operatorname{argmax}_{\alpha_i\in A_i} u_i(\alpha_i, a^*_{-i},\mu)$ and for any $a''_i \in \mathcal{A}_i$ satisfying $a''_i \neq a'_i$ there holds, $|u_i(a'_i, \hat{f}^i_{-i,\tau},\mu) - u_i(a'_i, a^*_{-i},\mu)| < \frac{\epsilon}{4}$ and $|u_i(a''_i, \hat{f}^i_{-i,\tau},\mu) - u_i(a''_i, a^*_{-i},\mu)| < \frac{\epsilon}{4}$. Adding and subtracting $u_i(a''_i, \hat{f}^i_{-i,\tau},\mu)$ to the right hand side of (13) and adding and subtracting $u_i(a'_i, \hat{f}^i_{-i,\tau},\mu)$ to the left hand side of (13), the preceding two inequalities give

$$u_i(a'_i, \hat{f}^i_{-i,\tau},\mu) - \frac{\epsilon}{2} > u_i(a''_i, \hat{f}^i_{-i,\tau},\mu).$$

By Assumption 1 and Lemma 3, $|u_i(a'_i, \hat{f}^i_{-i,\tau}, \mu^i_\tau) - u_i(a'_i, \hat{f}^i_{-i,\tau},\mu)| < \frac{\epsilon}{4}$ and $|u_i(a''_i, \hat{f}^i_{-i,\tau},\mu) - u_i(a''_i, \hat{f}^i_{-i,\tau}, \mu^i_\tau)| < \frac{\epsilon}{4}$ for all $t \geq t_\eta$; adding and subtracting these terms above gives

$$u_i(a'_i, \hat{f}^i_{-i,\tau}, \mu^i_\tau) > u_i(a''_i, \hat{f}^i_{-i,\tau}, \mu^i_\tau).$$

Since $\{a'_i\} = \operatorname{argmax}_{\alpha_i\in A_i} u_i(\alpha_i, a^*_{-i},\mu)$ and $a''_i \neq a'_i$ is arbitrary, this implies that $\operatorname{argmax}_{\alpha_i\in A_i} u_i(\alpha_i, \hat{f}^i_{-i,\tau}, \mu^i_\tau) = \operatorname{argmax}_{\alpha_i\in A_i} u_i(\alpha_i, a^*_{-i},\mu)$ for all i. □

A. Distributed Averaging in Dynamic Networks

Consider a network of n nodes connected through a communication graph $G = (V, E)$. The graph is assumed to be strongly connected. For $t = 1, 2, \ldots$ let $x_1(t) \in \mathbb{R}$ denote a value held by node 1 at time t. The objective is for all nodes to track as closely as possible the value $x_1(t)$.[5] Let $\epsilon(t) := |x_1(t+1) - x_1(t)|$ and assume that:

Assumption 7 *There exists a $B > 0$ such that $\epsilon(t) < B$ for all $t \in \mathbb{N}$.*

Let $\hat{x}_i(t)$ be the estimate player i maintains of $x_1(t)$. We make the following assumption pertaining to the initial error in players' estimates:

Assumption 8 $\hat{x}_i(0) - x_1(0) = 0$ *for all i.*

Let $\hat{x}(t) = (\hat{x}_1(t), \ldots, \hat{x}_n(t)) \in \mathbb{R}^n$ be a vector stacking all players' estimates, with $\hat{x}_1(t) = x_1(t)$ (i.e., node j knows its own value). Suppose the estimates are updated according to the following recursion:

$$\hat{x}(t+1) = W\left(\hat{x}(t) + e_1(x_1(t+1) - x_1(t))\right), \quad (14)$$

[5] In general, the objective may be to track the value $x_j(t)$ held by an arbitrary node j. Here, we only consider tracking $x_1(t)$, however, the general case is recovered by a permutation of the node labels.

where $e_1 \in \mathbb{R}^n$ is the 1-st canonical vector, and where the matrix $W \in \mathbb{R}^{n\times n}$ satisfies

Assumption 9 *W is row stochastic with sparsity conforming to G. Furthermore, W may be decomposed as*

$$W = \begin{pmatrix} 1 & 0 \\ b & P \end{pmatrix}$$

where $b \in \mathbb{R}^{n-1}$, $b \neq 0$ and P is irreducible (cf. [27]).

Note that in since W is row stochastic, $b \neq 0$ if and only if P is substochastic in the sense that at least one row sum of P is less than 1.

The following lemma gives a bound on the error in the agents' estimates of $x_1(t)$.

Lemma 10 *Suppose Assumptions 7–8 hold, and let the sequence $\{\hat{x}(t)\}_{t=1}^{\infty}$ be computed according to* (14)*. Suppose there exists a $t^* \geq 1$ and $T \geq 1$ such that $\sum_{t=t^*}^{t^*+T-2} \epsilon(t) \leq 1$ and $\{\epsilon(t)\}_{t=t^*}^{t^*+T-2}$ is decreasing. Then the error at time t^*+T-1 is bounded as,*

$$\|\hat{x}(t^*+T-1) - x_1(t^*+T-1)\mathbf{1}\| \leq \frac{n+1}{1-\lambda}\left(\frac{1}{T} + B\lambda^T\right),$$

where $\lambda := \sup_{\|y\|=1} \|Py\| < 1$.

Proof: Let $y(t) := \hat{x}(t) - x_1(t)\mathbf{1}$. Let $\delta(t) := [x_1(t+1) - x_1(t)]e_1 - [x_1(t+1) - x_1(t)]\mathbf{1}$. Subtracting $x_1(t+1)\mathbf{1}$ from both sides of (14) we get

$$\begin{aligned} y(t+1) &= W\left(\hat{x}(t) + [x_1(t+1) - x_1(t)]e_1\right) - x_1(t+1)\mathbf{1} \\ &= W\left(\hat{x}(t) + [x_1(t+1) - x_1(t)]e_1 - x_1(t+1)\mathbf{1}\right) \end{aligned}$$

where, in the second line, we may bring $x_1(t+1)\mathbf{1}$ inside the matrix multiplication due to the row stochasticity of W. Now we add and subtract $x_1(t)$ and use the definitions of $y(t)$ and $\delta(t)$ to get $y(t+1) = W\left(y(t) + \delta(t)\right)$. Inductively, this gives $y(t+1) = \sum_{s=0}^{t} W^{s+1}\delta(t-s) + W^{t+1}y(0)$. By Assumption 8 we have $y(0) = 0$, and hence $y(t+1) = \sum_{s=0}^{t} W^{s+1}\delta(t-s)$. By the triangle inequality we have

$$\|y(t+1)\| \leq \sum_{s=0}^{t} \|W^{s+1}\delta(t-s)\|. \tag{15}$$

Again using the triangle inequality, we establish a bound on $\|\delta(t)\|$:

$$\begin{aligned} \|\delta(t)\| &\leq \|[x_1(t+1) - x_1(t)]e_1\| + \|[x_1(t+1) - x_1(t)]\mathbf{1}\| \\ &\leq \epsilon(t) + n\epsilon(t) = (n+1)\epsilon(t). \end{aligned} \tag{16}$$

Let $\overline{W} := W - \mathbf{1}e_1^T$. In block form we have $\overline{W} = [0 \ \dots \ 0; \ (b - \mathbf{1}) \ \ P]$. Due to the special block form of $\overline{W}$, the spectrum $\sigma(\overline{W})$[6] of $\overline{W}$ consists precisely of $\{\sigma(P) \cup \{0\}\}$. Hence, the spectral radius of $\overline{W}$ coincides with that of P. In particular, the spectral radius of $\overline{W}$ is given by λ. Since P is substochastic, we have $\lambda < 1$.

Substituting $W = \overline{W} + \mathbf{1}e_1^T$ in (15) gives $\|y(t+1)\| \leq \sum_{s=1}^{t} \|(\overline{W} + \mathbf{1}e_1^T)^{s+1}\delta(t-s)\|$. Since $\overline{W}\mathbf{1} = \mathbf{0}$, and $e_1^T\overline{W} = \mathbf{0}$, and $\mathbf{1}e_1^T = (\mathbf{1}e_1^T)^s$ for any $s = 1, 2, \dots$, an inductive argument shows that $W^s = \overline{W}^s + \mathbf{1}e_1^T$ for any $s = 1, 2, \dots$. Thus we can upper bound $\|y(t+1)\|$ using the triangle inequality as follows

$$\|y(t+1)\| \leq \sum_{s=0}^{t} \left(\|\overline{W}^{s+1}\delta(t-s)\| + \|(\mathbf{1}e_1^T)^{s+1}\delta(t-s)\|\right)$$

It is readily verified that for $s = 1, 2, \dots$ there holds $e_1^T\delta(s) = 0$. Thus, the second term on the right hand side above is zero, i.e., $\|y(t+1)\| \leq \sum_{s=0}^{t} \|\overline{W}^{s+1}\delta(t-s)\|$. As a result we have, $\|y(t+1)\| \leq \sum_{s=0}^{t} \lambda^{s+1}\|\delta(t-s)\|$.

Let $t = t^* + T - 1$. Using the bound in (16) gives,

$$\begin{aligned} \|y(t)\| &\leq \sum_{s=0}^{t-1} \lambda^{s+1}(n+1)\epsilon(t-1-s) \\ &= (n+1)\sum_{s=0}^{t-t^*-1} \lambda^{s+1}\epsilon(t-1-s) + (n+1)\sum_{s=t-t^*}^{t-1} \lambda^{s+1}\epsilon(t-1-s). \end{aligned} \tag{17}$$

Consider the first term on the right hand side (RHS) above. Let $\epsilon_{avg}(t^*, T) := \frac{1}{T}\sum_{s=t^*}^{t^*+T-2} \epsilon(s)$. By assumption, the sequence $\{\epsilon(t)\}_{t=t^*}^{t^*+T-2}$ is decreasing, hence by Chebychev's sum inequality [28] (p. 43-44), $\sum_{s=0}^{t-t^*-1} \lambda^{s+1}\epsilon(t-1-s) \leq \epsilon_{avg}(t^*, T)\sum_{s=0}^{t-t^*-1} \lambda^{s+1} \leq \epsilon_{avg}(t^*, T)\frac{1}{1-\lambda}$, where the latter inequality follows by taking the closed form of the geometric sum. Furthermore, by assumption we have $\sum_{s=t^*}^{t^*+T-2} \epsilon(s) \leq 1$, and hence $\epsilon_{avg}(t^*, T) \leq \frac{1}{T}$, which gives that $\sum_{s=0}^{t-t^*-1} \lambda^{s+1}\epsilon(t-s) \leq \frac{1}{T}\frac{1}{1-\lambda}$.

Now consider the second term on the RHS of (17). By Assumption 7, we have $\epsilon(t-s) \leq B$ which allows us to bound the second term as $(n+1)\sum_{s=t-t^*}^{t-1} \lambda^{s+1}\epsilon(t-1-s) \leq (n+1)B\sum_{s=t-t^*}^{t-1} \lambda^{s+1} = (n+1)B\lambda^T\sum_{s=0}^{t^*-1} \lambda^s \leq (n+1)B\lambda^T(1-\lambda)^{-1}$, where the latter inequality again follows by taking the closed form of the geometric sum.

Substituting these two bounds back into (17) we get

$$\begin{aligned} \|y(t)\| &\leq (n+1)\frac{1}{T}\frac{1}{1-\lambda} + (n+1)B\lambda^T\frac{1}{1-\lambda} \\ &= (n+1)\frac{1}{1-\lambda}\left(\frac{1}{T} + B\lambda^T\right). \end{aligned}$$

Since we chose $t = t^* + T - 1$, this concludes the proof. □

In order to apply Lemma 10, one must show that $\sum_{t=t^*}^{t^*+T-2} \epsilon(t) < 1$. Essentially, this condition states that the variation in the node value $x_1(t)$ during the designated time interval remains bounded. This can be easily ensured, for example, if the value of $x_1(t)$ is monotone. This is the content of the following Lemma.

Lemma 11 *Suppose that $x_1(t) \in [0, 1]$ for all t. Suppose also there exist $t^*, T \in \mathbb{N}_+$ such that $\{x_1(t)\}_{t=t^*}^{t^*+T-1}$ is a monotone sequence. Then $\sum_{t=t^*}^{t^*+T-2} \epsilon(t) \leq 1$.*

Proof: Suppose that $\{x_1(t)\}_{t=t^*}^{t^*+T-1}$ is monotone increasing. Then $\epsilon(t) = |x_1(t+1) - x_1(t)| = x_1(t+1) - x_1(t)$. Substituting this into the sum in question gives a telescoping sum $\sum_{t=t^*}^{t^*+T-2} \epsilon(t) = \sum_{t=t^*}^{t^*+T-2} x_1(t+1) - x_1(t) = x_1(t^*+T-1) - x_1(t^*) \leq 1$. The final inequality follows since $0 \leq x_1(t) \leq 1$ for all t. A similar argument handles the monotone decreasing case. □

We now prove Lemma 5 of Section IV-F.

Proof (Lemma 5): Let $\epsilon > 0$ and let $t^* \in \mathbb{N}_+$ be arbitrary. Our task is to show that under the update rule (9), there exists a T such that if starting at time t^* any action is repeated in $\tilde{T} \geq T$ consecutive stages, then $\|\hat{f}^i_{t^*+\tilde{T}-1} - f_{t^*+\tilde{T}-1}\| < \epsilon$. We will accomplish this by showing that the update rule (9) fits the template of Lemma 10.

Fix a player $j \in \mathcal{N}$ and action $a_j \in \mathcal{A}_j$. Let $f_{j,t}(a_j)$ denote the weight that the empirical distribution $f_{j,t}$ places on a_j, and similarly, let $\hat{f}^i_{j,t}(a_j)$ denote the weight that $\hat{f}^i_{j,t}$ places on a_j. For the purpose of applying Lemma 10, let $x_j(0) = 0$, let $x_j(t) := f_{j,t}(a_j)$, $t \geq 1$ and for $i = 1, \dots, n$ let $\hat{x}_i(0) = 0$, and let $\hat{x}_i(t) := \hat{f}^i_{j,t}(a_j)$ for $t \geq 1$. By (9) and the initialization condition for Algorithm 3 for $t \geq 0$ we have $\hat{x}_i(t+1) = \sum_{k\in\mathcal{N}_i} w^i_{j,k}\left(\hat{x}_k(t) + (x_j(t+1) - x_j(t))\chi_{\{k=j\}}\right)$. Letting $\hat{x}(t) =$

[6] In Section II the symbol σ was used to represent a mixed strategy. In keeping with standard conventions, we use σ here to denote the spectrum of a matrix, where the distinction is clear from the context.

$(\hat{x}_i(t))_{i=1}^n \in \mathbb{R}^n$ we may express the update rule in more compact notation as

$$\hat{x}(t+1) = W_j\left(\hat{x}(t) + e_j(x_j(t+1) - x_j(t))\right),$$

where $W_j = (w_{j,k}^i)_{i,k\in\mathcal{N}}$ is the weight matrix as assumed in Lemma 5 and e_j is the j-th canonnical vector in $\mathbb{R}^{|\mathcal{A}_j|}$. Note that, after a permutation of the player ordering (which causes no loss in generality), this fits the format of (14). Note also that Assumption 7 is satisfied since $x_j(t) \in [0,1]$ for all t. Furthermore, Assumption 8 is satisfied since, by construction, $\hat{x}_i(0) = x_j(0) = 0$ for all i, and Assumption 9 is satisfied by the hypothesis of Lemma 5.

Now, let $\epsilon(t) := |x_j(t+1) - x_j(t)|$ and suppose that starting at time t^* some action $a^* = (a_1^*, \ldots, a_n^*) \in \mathcal{A}$ is repeated in T consecutive stages, where $T \in \mathbb{N}_+$ is arbitrary. Two cases must be considered—the case that $a_j = a_j^*$ (i.e., the action which defines x_j is in fact the action being repeated by player j), and the case that $a_j \neq a_j^*$ (i.e., the action which defines x_j is not being played at all by j during the designated time sequence.) If $a_j = a_j^*$ then $\{x_j(t)\}_{t=t^*}^{t^*+T-1} = \{f_{j,t}(a_j)\}_{t=t^*}^{t^*+T-1}$ increases monotonically towards 1 (this follows from (8)). Otherwise, if $a_j \neq a_j^*$ then $\{x_j(t)\}_{t=t^*}^{t^*+T-1} = \{f_{j,t}(a_j)\}_{t=t^*}^{t^*+T-1}$ decreases monotonically towards 0. Since, in either case the sequence is monotone, we have by Lemma 11 that $\sum_{t=t^*+1}^{t^*+T-2} \epsilon(t) \leq 1$.

Note also that if some action a^* is repeated from time t^* to t^*+T-1, then the difference sequence $\{\epsilon(t)\}_{t=t^*}^{t^*+T-2} = \{|f_{j,t+1}(a_j) - f_{j,t}(a_j)|\}_{t=t^*}^{t^*+T-2}$ is decreasing. This follows from (8).

We are now in a position to apply Lemma 10. By the equivalence of finite dimensional norms, there exist constants c_1 and c_∞ such that $\|\cdot\| \leq c_1\|\cdot\|_1$ and $c_\infty\|\cdot\|_\infty \leq \|\cdot\|$. Given $j \in \mathcal{N}$ and some action $a_j \in \mathcal{A}_j$ we may choose a constant $T_{a_j} \in \mathbb{N}_+$ sufficiently large such that $\frac{n+1}{1-\lambda}\left(\frac{1}{T_{a_j}} + B\lambda^{T_{a_j}}\right) < c_\infty \frac{\epsilon}{c_1\sum_{i=1}^n |\mathcal{A}_i|}$. Applying Lemma 10 we get that if any action a^* is repeated in $T \geq T_{a_j}$ consecutive stages starting at any time t^* then

$$\begin{aligned}
&\max_{i\in\mathcal{N}} |f_{j,t^*+T-1}^i(a_j) - f_{j,t^*+T-1}(a_j)| \qquad (18)\\
&= \|(f_{j,t^*+T-1}^i(a_j))_{i=1}^n - f_{j,t^*+T-1}(a_j)\mathbf{1}\|_\infty\\
&= \|\hat{x}(t^*+T-1) - x_j(t^*+T-1)\mathbf{1}\|_\infty\\
&< \frac{1}{c_\infty}\|\hat{x}(t^*+T-1) - x_j(t^*+T-1)\mathbf{1}\| \leq \frac{\epsilon}{c_1\sum_{i=1}^n |\mathcal{A}_i|}.
\end{aligned}$$

Let $T := \max_{j\in\mathcal{N}, a_j\in\mathcal{A}_j} T_{a_j}$. By (18) we have $|f_{j,t^*+\tilde{T}}^i(a_j) - f_{j,t^*+\tilde{T}}(a_j)| < \frac{\epsilon}{c_1|\sum_{i=1}^n |\mathcal{A}_i|}$ for all $\tilde{T} \geq T$ for all $i,j \in \mathcal{N}$ and for all $a_j \in \mathcal{A}_j$.

Now, fix any player $i \in \mathcal{N}$. Observe that for any $\tilde{T} \geq T$ we have $\|\hat{f}_{t^*+\tilde{T}-1}^i - f_{t^*+\tilde{T}-1}\| \leq c_1\|\hat{f}_{t^*+\tilde{T}-1}^i - f_{t^*+\tilde{T}-1}\|_1 = c_1\sum_{j\in\mathcal{N}}\sum_{a_j\in\mathcal{A}_j} |\hat{f}_{j,t^*+\tilde{T}-1}^i(a_j) - f_{j,t^*+\tilde{T}-1}(a_j)| \leq c_1(\sum_{j=1}^n |\mathcal{A}_j|)\frac{\epsilon}{c_1\sum_{j=1}^n |\mathcal{A}_j|} = \epsilon$, which is the desired result. □

REFERENCES

[1] J. R. Marden and J. S. Shamma, "Game theory and distributed control," *Handbook of game theory*, vol. 4, pp. 861–899, 2012.

[2] S. Lasaulce, M. Debbah, and E. Altman, "Methodologies for analyzing equilibria in wireless games," *IEEE Signal Processing Magazine*, vol. 5, no. 26, pp. 41–52, 2009.

[3] G. Bacci, S. Lasaulce, W. Saad, and L. Sanguinetti, "Game theory for networks: A tutorial on game-theoretic tools for emerging signal processing applications," *IEEE Signal Processing Magazine*, vol. 33, no. 1, pp. 94–119, 2016.

[4] L. Rose, S. Lasaulce, S. M. Perlaza, and M. Debbah, "Learning equilibria with partial information in decentralized wireless networks," *IEEE Communications Magazine*, vol. 49, no. 8, pp. 136–142, 2011.

[5] W. Saad, Z. Han, H. V. Poor, and T. Basar, "Smart meters for power grid: Challenges, issues, advantages and status," *IEEE Signal Process. Mag.*, vol. 29, no. 5, pp. 86–105, 2012.

[6] G. Arslan and J. Shamma, "Distributed convergence to nash equilibria with local utility measurements," in *43rd IEEE Conference on Decision and Control*, vol. 2, 2004, pp. 1538–1543.

[7] D. Fudenberg and D. Levine, *The Theory of Learning in Games*, 1st ed. Cambridge, MA: MIT Press, 1998.

[8] G. Arslan, J. Marden, and J. Shamma, "Autonomous vehicle-target assignment: a game-theoretical formulation," *Journal of Dynamic Systems, Measurement, and Control*, vol. 129, no. 5, pp. 584–596, 2007.

[9] J. Marden, G. Arslan, and J. Shamma, "Joint strategy fictitious play with inertia for potential games," *IEEE Trans. Automatic Control*, vol. 54, no. 2, pp. 208–220, 2009.

[10] H. Young, *Strategic learning and its limits*. Oxford University Press, 2004.

[11] J. R. Marden, S. D. Ruben, and L. Y. Pao, "A model-free approach to wind farm control using game theoretic methods," *IEEE Tran. on Control Systems Technology*, vol. 21, no. 4, pp. 1207–1214, 2013.

[12] B. Swenson, S. Kar, and J. Xavier, "Distributed learning in large-scale multi-agent games: A modified fictitious play approach," in *Signals, Systems and Computers (ASILOMAR), 2012 Conference Record of the Forty Sixth Asilomar Conference on*. IEEE, 2012, pp. 1490–1495.

[13] F. Salehisadaghiani and L. Pavel, "Nash equilibrium seeking by a gossip-based algorithm," in *Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on*. IEEE, 2014, pp. 1155–1160.

[14] J. Koshal, A. Nedic, and U. V. Shanbhag, "A gossip algorithm for aggregative games on graphs," in *IEEE 51st Annual Conference on Decision and Control (CDC)*, 2012, pp. 4840–4845.

[15] B. Swenson, S. Kar, and J. Xavier, "Empirical centroid fictitious play: An approach for distributed learning in multi-agent games," *IEEE Trans. Signal Process.*, vol. 63, no. 15, pp. 3888 – 3901, 2015.

[16] C. Eksin and A. Ribeiro, "Distributed fictitious play in potential games of incomplete information," in *IEEE 54th Annual Conference on Decision and Control (CDC)*, Osaka, Japan, 2015, pp. 5190 – 5196.

[17] G. W. Brown, "Iterative solution of games by fictitious play," *Activity analysis of production and allocation*, vol. 13, no. 1, pp. 374–376, 1951.

[18] J. Shamma and G. Arslan, "Dynamic fictitious play, dynamic gradient play, and distributed convergence to nash equilibria," *IEEE Trans. Automatic Control*, vol. 50, no. 3, pp. 312–327, 2005.

[19] A. Fabrikant, A. D. Jaggard, and M. Schapira, "On the structure of weakly acyclic games," in *Algorithmic Game Theory*. Springer, 2010, pp. 126–137.

[20] D. Williams, *Probability with martingales*. Cambridge university press, 1991.

[21] R. Olfati-Saber and J. S. Shamma, "Consensus filters for sensor networks and distributed sensor fusion," in *Decision and Control, 2005 and 2005 European Control Conference. CDC-ECC'05. 44th IEEE Conference on*. IEEE, 2005, pp. 6698–6703.

[22] R. Carli, A. Chiuso, L. Schenato, and S. Zampieri, "Distributed kalman filtering based on consensus strategies," *Selected Areas in Communications, IEEE Journal on*, vol. 26, no. 4, pp. 622–633, 2008.

[23] N. Xiao, X. Wang, T. Wongpiromsarn, K. You, L. Xie, E. Frazzoli, and D. Rus, "Average strategy fictitious play with application to road pricing," in *2013 American Control Conference*. IEEE, 2013, pp. 1920–1925.

[24] A. G. Dimakis, S. Kar, J. M. Moura, M. G. Rabbat, and A. Scaglione, "Gossip algorithms for distributed signal processing," *Proceedings of the IEEE*, vol. 98, no. 11, pp. 1847–1864, 2010.

[25] S. Kar and J. M. Moura, "Distributed consensus algorithms in sensor networks with imperfect communication: Link failures and channel noise," *Signal Processing, IEEE Transactions on*, vol. 57, no. 1, pp. 355–369, 2009.

[26] A. H. Sayed, "Diffusion adaptation over networks," *Academic Press Library in Signal Processing*, vol. 3, pp. 323–454, 2013.

[27] U. A. Khan, S. Kar, and J. M. Moura, "Higher dimensional consensus: Learning in large-scale networks," *IEEE Transactions on Signal Processing*, vol. 58, no. 5, pp. 2836–2849, 2010.

[28] G. H. Hardy, J. E. Littlewood, and G. Polya, *Inequalities, Cambridge Mathematical Library*. Cambridge University Press, 1988.